PRTN # 253
Stan Lieberson
Fouad Tobagi
UCLA

CURRENT PR MEASUREMENT PLAN, MARCH, 1978

This plan constitutes the description of a number of measurement activities to be conducted on the experimental PRNET. It is expected and intended that this document will be modified and expanded to reflect increasing knowledge of the capabilities, limitations and parameters of the experimental PRNET, as well as changes in PRNET protocol and hardware.

The plan consists of the following:

Part I: Objectives of the measurement program including a description of the PRNET measurement facilities.

Part II: A list of specific network experiments, along with the required configurations, tools, parameters and data reduction.

Part III: Actual test bed configurations to carry out the network experiments listed in Part II. (Drawings obtained from SRI.)

Part IV: A tentative time table for the execution, data reduction and analysis of the experiments in Part II.

PART I

MOTIVATION AND OBJECTIVES OF THE PACKET RADIO MEASUREMENT PROGRAM

D.O.D PACKET RADIO COMMUNICATIONS REQUIREMENTS

The advantages in using broadcast radio communications are many: easy access to central computer installations and computer networks; collection and dissemination of data over local distributed geographical areas independent of the availability of preexisting (telephone) wire networks; the suitability of wireless connections for comunications with and among mobile users.

The military applications are further supported by the potential rapid deployment of packet radio systems, the mixing of voice and data transmission over the same communications media and the achievement of distributed control of shared channels rendering the system both cost-effective and reliable.

THE NEED FOR THE PACKET RADIO MEASUREMENT EFFORT

The concern for measurement is due to several factors. Firstly, these measurements provide a means to evaluate the performance of the protocols employed and the identification of their key parameters. Moreover, this realistic observation of the system behavior will assist in the validation and improvement of existing analytical models devised to study some of these schemes, such as the access modes and routing strategies. Secondly, these measurements will allow for the detection of system inefficiencies and the identification of design flaws such as the inadvertent creation of a deadlock condition. Thirdly, measurement facilities and data, when used to improve network design, are a valuable feedback process in which design deficiencies are detected and subsequently corrected.

MEASUREMENT FACILITIES

Immediately below we describe the various types of statistics to be collected in the Packet Radio Net, the traffic sources required in measurement experiments and the techniques available for measurement data collection.

CUMULATIVE STATISTICS (CumStats)

As their name suggests these consist of data regarding a variety of events, accumulated over a given period of time, and proveded in the form of sums, frequencies and histograms. We shall distinguish between those data collected at the PRUs (PRU based CumStats) and those collected at the end devices (the end-to-end CumStats). The PRU based CumStats provide information about the local environment and behavior such as traffic load, channel access, routing performance, and repeater activity. Conversely, end-to-end statistics collected at network sources and sinks, that is stations and terminal devices, will reflect more global network behavior such as user delays and network throughput.

TRACE STATISTICS

The trace capability allows one to literally follow a packet through the network, and to trace the route which it takes and the delays which it encounters at each hop. In the packet radio network, the collection of trace data to the repeaters is prohibited by the limited size of storage in the PRU. To overcome this problem, we have introduced a new type of packet called the Pickup Packet. These packets are generated with an empty text field by traffic generators at end devices. As these packets flow normally in the network according to the transport protocols, selected repeaters will gather the trace statistics and will store them within the text field of the pickup packets themselves.

SNAPSHOT STATISTICS

Snapshots give an instantaneous peek at a PRU, showing its state at that moment with regard to buffer assignment and queue lengths. (In the ARPANET, which is a decentralized network in which each node contains routing algorithms and data, snapshots also include routing related information; in the Packet Radio Network, such information is available at the station). Changes to appropriate station tables will be time stamped and collected as the station's snapshot function.

ARTIFICIAL TRAFFIC GENERATORS

The creation of streams of packets between given points in the net, with given durations, intervals, packet lengths, and

packet types (Information and Pickup Packets) is clearly a requirement of any experimental system. While it might be desirable to provide each PRU with the capability of creating such traffic, this additional burden on the PRU software can be avoided if there exists a reasonable number of terminals with processors attached which, along with the station, will be programmed to provide the traffic-source functions indicated above.

Specifically, traffic-source features which the terminals (and Station) should provide are:

(1) Information Packets - the user specifies the packet length, frequency, destination and duration of one or more streams of Information Packets. (The text content may be arbitrary.)

(2) Pickup Packets - the user specifies the packet length, frequency, destination and duration of one or more streams of Pickup Packets.

STATION MEASUREMENT PROCESS

Since the station is the central node and provides central control for the operation of the entire network it therefore plays a central role in the execution of measurement functions. It is through the station that the initiation and termination of measurement experiments are controlled. In particular, the station enables and disables the CumStat and Pickup packet functions at the PRU's, and assigns to the various elements the intervals for CumStat collections, and to the artificial traffic generator, their corresponding parameters. Moreover, it is to the station that all measurement data is ultimately destined; upon arrival at the station, the data is time-stamped and stored in a single measurement file for off-line reduction and analysis. In addition, all changes to the station's internal tables (routing, connectivity, PRU operational parameters, etc.) will be reflected by an entry into the measurement file, thus allowing the correlation of measurement results to the actual network configuration. A (Measurement) process at the station will perform all of the above functions.

SUMMARY OF MEASUREMENT ITEMS

Pickup Packets (at each PRU, the following data items are collected in the Pickup packet):

Time of arrival of the packet at the PRU
Time the Pickup packet was first placed on the transmit queue
Time of each transmission

Time HBH ack arrived (stored in next Pickup packet)
The current PRU ID

PRU based Cumulative Statistics

Number of packets received in error
Number of packets received but not intended for this PRU
Histogram of number of transmissions per successful packet
Number of unsuccessful packets (dropped because of lack of ack)
Number of packets discarded because of lack of buffer space
Number of alternately routed ("ALL") packets received
Table counting number of correctly received packets from immediate neighbors
Number of transmissions beyond success
Table sampling frequency of buffer states (and transceiver states)

End-Device Cumulative Statistics

Histogram of round-trip times
Number of packets transmitted
Number of duplicate packets detected
Number of packets discarded by the sender because of lack of ETE ack
Histogram of number of transmissions per successful (ETE) packet
Histogram of packet intergeneration time

Note: certain CumStat items will distinguish between inbound (to the station) and outbound (from the station) traffic, as well as 100 and 400 KBPS Rates.

PART II

PROPOSED PACKET RADIO MEASUREMENTS

Table of Symbols:

- ○ PRU
- □ Station
- △ Terminal Device
- ---- Routing Assigned
- ⬭ Area of Bi-Directional Connectivity

EVALUATE CHANNEL ACCESS PROTOCOLS IN ONE-HOP ENVIRONMENTS

In the initial experimental system pure ALOHA, "disciplined ALOHA" and non-persistent CSMA will be available. Our measurement aims are to validate the analytical models of the three access modes and to evaluate their performance in realistic environments.

In evaluating terminal access in a single hop system (a model commonly used in analysis), we consider an environment consisting of a single station and a population of terminals within range and in line-of-sight of the station. In order to determine the relationships between the network throughput (rate of successfully received packets at the station) and channel traffic (rate of packet transmissions over the channel), as well as the relationships between the network throughput and packet delays, the following quantities will be measured:

(a) the number of transmissions a packet incurs before success

(b) the one-hop packet delay: time elapsed since the packet is ready for transmission until it is acknowledged, i.e., until its acknowledgment packet is received from the station

(c) network throughput: average number of packets received at the station per unit time

Items (a) and (b) are obtained in the form of histograms by the CumStat tools at the PRU and the end device respectively. Item (b) may also be obtained individually for each Pickup packet by having the originating device store the time its acknowledgment arrived. Item (c) is obtained at the station from end-to-end cumulative statistics.

Note: When comparing measurement results with analysis, it is important to note the following important discrepency: the analysis assumed that acknowledgement packets are instantaneous and for free; the measurement results will bear the effects of acknowledgement traffic, of its interference with information packets and of the loss of acks.

Area of Interest: 1-hop Channel Access

See SRI testbed figure(s): 2 (modified by moving a terminal)

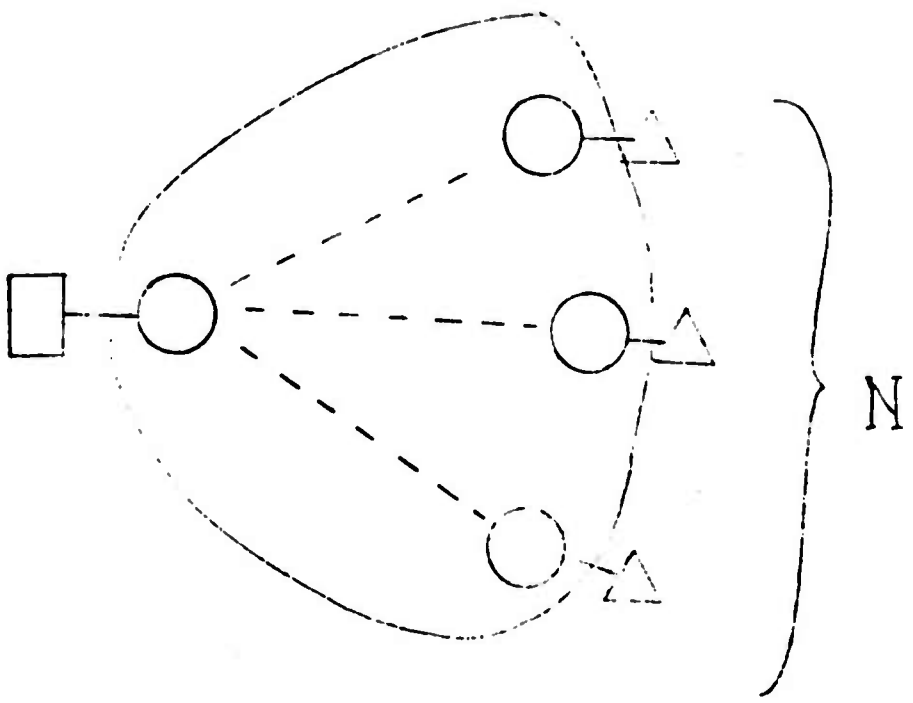

Area of Interest: 1-hop Channel Access

Configuration: A 1-hop network with all involved PRUs in line-of-sight and within range

Tools: Cumstats (End-device and PRU), Pickup Packets

Data Items:
1) # of tx's to successful acknowledgement [CS-PRU]
2) # of successful tx's [CS-PRU] (number of tx's HBH ack'ed)
3) 1-hop (one-way) delay [Pickup Packets]
4) # of packets successfully rx'd at the station [CS-station]
5) # of packets successfully rx'd at the receiving PRU (station's PRU) [CS-PRU]

Parameters: Traffic sources sending information packets (one run with minimum size packets, one run with maximum size packets) sent at a range of intergeneration times, the range to Initial experiments will determine the limits of this range.

Retransmission delays (mean) will vary within a range of values over several runs of this experiment. (Range will be determined experimentally.)

Processing: mean, SD, pdf of: # of tx's, plot of:

a) delay vs throughput

b) # of tx's to success vs throughput

c) maximum throughput vs N

d) maximum throughput vs retransmission delay

Comments: Each of the N configurations to be run for the two packet sizes, a range of traffic rates and mean retransmission times, and compared across the three access methods (ALOHA, "disciplined ALOHA", non-persistent CSMA). Other parameters to note: i) effect of signal power (capture), ii) all 400KB or 100KB rates.

EVALUATE CHANNEL ACCESS PROTOCOLS IN MULTI-REPEATER ENVIRONMENTS

The tasks of measuring performance of terminal access techniques in multi-repeater environments differs from the previous one in that repeater-to-repeater traffic is present contending on the same channel. The environment consists of a number of repeaters and stations and a population of terminals, not necessarily all within range and in line-of-sight. The same quantities as listed above, measured over the terminal-to-repeater hop, will be collected using the same tools.

Area of Interest: Multi-Repeater 1-hop Channel Access

See SRI testbed figure(s): 1

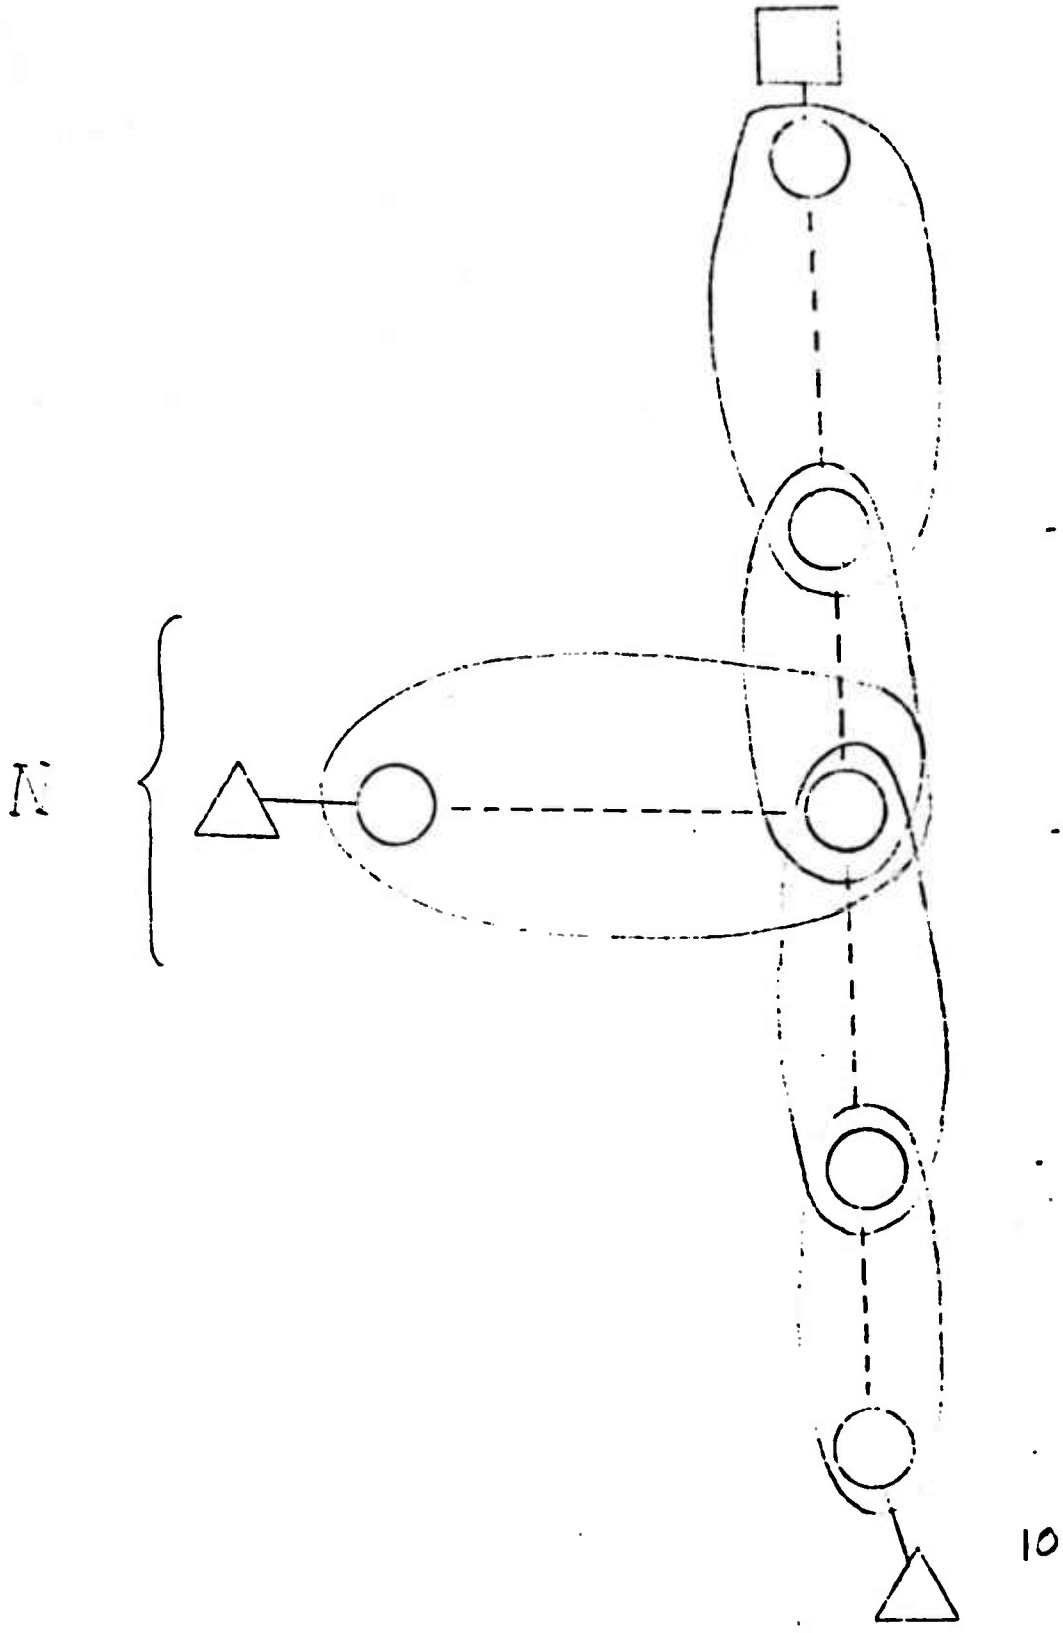

Area of Interest: Multi-Repeater 1-hop Channel Access

Configuration: A pool of terminals accessing a repeater, with contending network traffic routed through that repeater.

Tools: Cumstats (end-device and PRU), Pickup packets

Parameters: The N terminals send information packets. The comments as to packet sizes and intergeneration times for 1-hop channel access apply here. The end terminal will provide a stream of information packets to the station at a similiar range of rates.

Data Items:
1) # of tx's to success [CS-PRU]
2) 1-hop (first hop) delay [PP]
3) # of packets successfully received at the station [CS-station]
4) # of successful tx's [CS-PRU]

Processing: Same as for 1-hop channel access, but with all results plotted against level of interference caused by transit traffic (end terminal data rate), and N.

HIDDEN TERMINAL

In these one-hop channel access experiments, the performance of CSMA is dependent of the connectivity pattern among the contending terminals. The "hidden terminal" problem can, as shown by analysis, significantly degrade such a permormance. This experiment is aimed at a characterization of the effect of the hidden terminal problem on the performance measures discussed in the preceeding experiments (One-Hop Channel Access; Multi-Repeater One-Hop Channel Access).

Area of Interest: Hidden Terminal Problem in Non-Persistent CSMA

See SRI testbed figure(s): 2

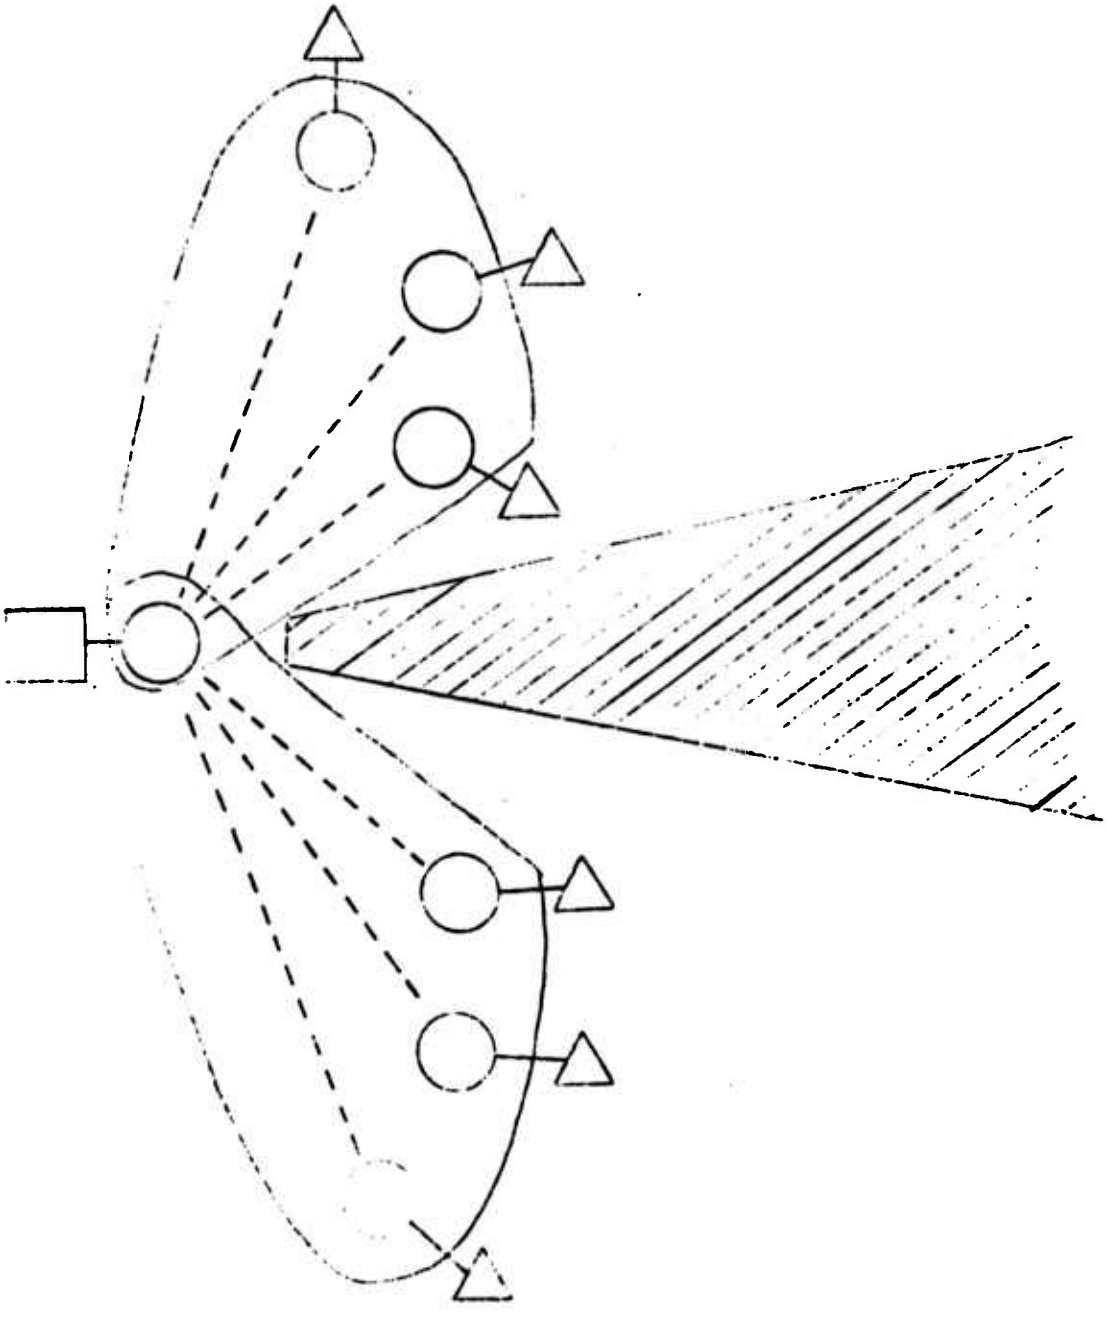

Area of Interest: Hidden Terminal Problem in Non-Persistent CSMA

Configuration: Two groups of terminals, with no connectivity between the groups, accessing the same PRU.

Tools: Cumstats (end-device and PRU)

Parameters: Each group sending α and $1-\alpha$ of the traffic, respectively.

Data Items:
1) # of tx's to success [CS-PRU]
2) 1-hop delay [CS-terminal]
3) # of packets successfully rx'd (ETE) [CS-station]
4) # of successful transmissions

Processing: [same as in previous two channel-access experiments] Compare with analytical results of hidden terminal problem.

HOP-BY-HOP ACKNOWLEDGMENT PROTOCOL

Hop-by-hop acknowledgment suffers from packet interference. The delay until the echo acknowledgment is received at the transmitter (provided the ack is not lost or conflicted with) is random. Thus, the packet may incur some additional transmissions beyond the first successful one creating additional overhead on the channel and the PRU's. This number of additional transmissions is a measure of the inefficiency of echo acknowledgments; so too will be the number of packets discarded at the transmitter because of lack of reception of the echo acknowledgment. That is, the transmitter reached the maximum number of retransmissions of a packet before the echo acknowledgment was received; although the packet may have been successful, the transmitter declares itself unsuccessful in establishing communication!

Thus, we shall measure the efficiency of the Echo Acknowledgment protocol by measuring the number of additional transmissions beyond success incurred by a packet. To compute this number, a PRU must have two pieces of information; it must know how many times the packet has been transmitted, and it must also know which of those retransmissions was the one that reached the next repeater successfully. This information will be contained in two fields in each packet header, which we refer to here as fields A and B. Field B is used by the PRU to store the current transmission number of the packet. When the packet is successfully heard by the intended receiver, the contents of field B are saved by being stored into field A; when the Echo acknowledgment is successfully heard by the sending PRU, field A of the echo acknowledgment is compared with the current number of transmissions of the packet, i.e., the contents of field B in the sender's copy of the packet. If these two numbers differ, then the magnitude of that difference represents the number of times that the packet was retransmitted after it had already been successfully received at the next hop. This data is collected as part of the cumulative statistics of the sending PRU.

The independent parameters for these experiments are the input traffic rates, the retx delay, and the time-out periods at each device. The data items to be collected are the number of tx's beyond success as well as packet delays (either end-to-end or single hop delays). An increase in the "time-out" period may decrease the number of tx's beyond success, but they may increase the delay. [The time-out period is that minimum delay each packet must experience before it can be transmitted, i.e. the fixed time to which the random delay is added.]

Area of Interest: Hop-by-hop acknowledgement protocols

See SRI testbed figure(s): 1

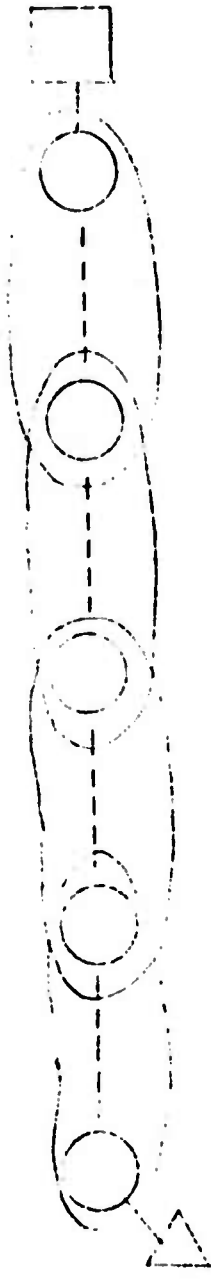

Area of Interest: Hop-by-hop acknowledgement protocols

Configurations: two basic configurations will be considerd: a string of repeaters of some length; and a ring of terminals around a station, forming a single hop system.

Tools: Cumstats (PRU)

Parameters: For both configurations, we vary the input rates at the traffic sources, the retx delays and time-out periods for rescheduling.

Data Items: 1) # of tx's beyond success.
2) packet delays, either end-to-end (round-trip) or single hop delays.

Processing: mean, SD, pdf of: # of tx's beyond succcess (vs N, and vs traffic load), plot: # of tx's beyond success vs traffic load (for N)

Evaluation of Repeater's Performance:

The evaluation of the performance of a repeater is most important in the anaylsis of network behavior; it allows us to break down key network measures (such as packet delay and throughput) into their elementary components and to examine the effects on these measures of the repeater activity and design (including buffer management, queueing discipline, and packet processing priorities).

PACKET PROCESSING DELAYS

The quantities relevant to packet delay at a PRU are:

(a) The processing time of a packet flowing through a repeater; this is counted in Pickup Packets as the time lapse between the packet's arrival and the time it is placed on the transmission queue. This processing includes various checks such as checksum, packet type, routing tables. etc.

(b) the packet queueing delay at a repeater; this is also counted in Pickup Packets as the time elapsed from when the packet is placed on the transmission queue until it is considered for transmission (i.e., until it is at the head of the line, in a first-come-first-served discipline).

(c) the packet's service time; this is also counted in Pickup packets as the time elapsed from when the packet is at the head of the queue until its HBH acknowledgment is correctly received. Note that the actual service time (time until the packet is correctly received at the next repeater) is smaller than the one measured here due to the echo acknowledgment protocol used in this system. Note also that the service times of consecutive packets may be correlated.

Area of Interest: Packet processing time

See SRI testbed figure(s): 1

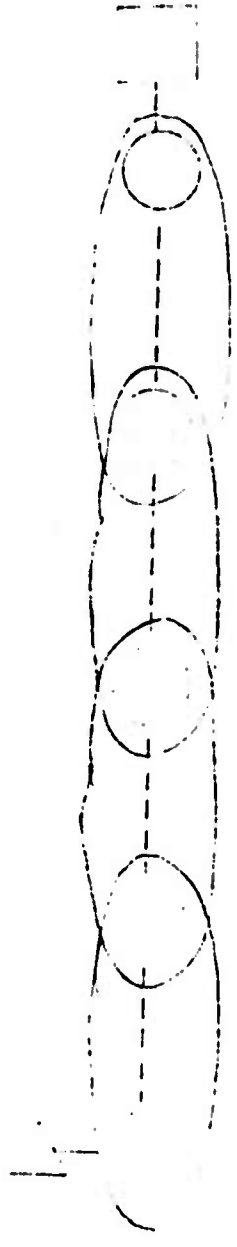

Area of Interest: Packet processing time

Configurations: A string of PRUs appears to be the most obvious configuration to consider. A minimum of 5 PRUs guarantees the existence of all combinations of data rates (100/400 KBPS) and HBH ack types (echo/active).

Tools: Pickup Packets

Parameters: Varying input rates at traffic source, varying retx delays at PRUs.

Data Items: (as defined in text above)
Processing time (PP)
Queueing delay (PP)
Service time (PP)

Processing: For each of the PRU's, plot -
processing time, queueing delay and service time vs traffic load and retx delay.

TRANSCEIVER ACTIVITY

The quantities related to a repeater's communications activity are:

(a) percent of time the PRU transceiver is busy transmitting and receiving; this can be obtained in the PRU CumStat by regular sampling of the transceiver's state.

(b) the percent of traffic received with checksum error (obtained in the PRU CumStats).

(c) the percent of traffic received correctly but not intended for this repeater (obtained in the PRU CumStats).

Area of Interest: Transceiver activity

See SRI testbed figure(s): b) 1; c) 5,6 [modified]

a)

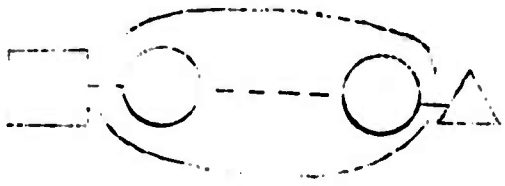

b)

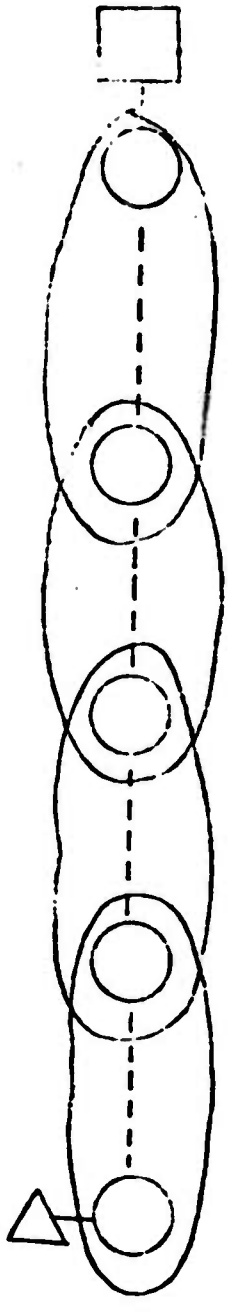

c)

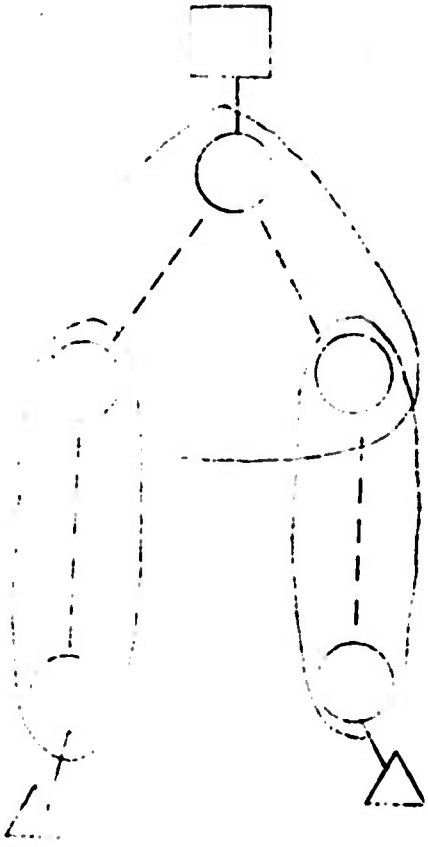

Area of Interest: Transceiver activity

Configurations: (a) for the simplest configuration; (b) for a string containing all combinations of data rates and HBH acks (100/400 and active/echo); (c) for more generality in contention among PRUs

Tools: Cumstats (PRU), Snapshots (PRU)

Parameters: various input rates, retx delays.

Data Items:

Transceiver States:

fraction of time transceiver is:

a) disabled
b) tx'ing
c) rx'ing
d) rx enabled (listening)

Other Items:

- # of words tx'd
- # of packets correctly tx'd
- # of packets correctly rx'd
- # of packets incorrectly rx'd

Processing: plot of: frequency of above transceiver states vs traffic load, fraction of correct tx's, fraction of correct rx's.

Comments: The size and timing of the transceiver disabled state ("black hole") is of special concern.

BUFFER MANAGEMENT

The quantities relevant to buffer management and occupancy are:

(a) the percent of time packet buffers are in a given state (free, queued for packet transmission, reserved for packet receive). This can be obtained in the PRU CumStat by a regular sampling of the buffer states.

(b) the frequency of buffer overflow as a function of the load; this is obtained also in the CumStats by counting the number of packets discarded due to lack of buffer space.

Area of Interest: PRU buffer Management

Configurations: (a) for the simplest configuration; (b) for a string containing all combinations of data rates and HBH acks (100/400 and active/echo); (c) for more generality in contention among PRUs [See Transceiver Activity]

Tools: Cumstats (PRU), Snapshots (PRU)

Parameters: various input rates, retx delays.

Data Items:
Buffer States:
fraction of time buffers are:
a) free
b) queued for packet tx
c) reserved for packet rx (radio 100 KB, 400 KB; sta/ter)

Other Items:
- number of packets correctly tx'd
- number of packets correctly rx'd
- number of packets incorrectly rx'd

Processing: plot of: frequency of above buffer states vs traffic load

Routing Protocols:

In this experimental network, the hierarchical routing scheme(*) in use is based on a tree structure with the station as its root. The initial tree structure is created dynamically by the Initialization Procedure in which the station used PRU connectivity information to create a tree that attempts to "minimize" the number of hops between each repeater and the station. (However, when the first choice shortest path cannot be used, the packet departs from this path and uses an alternate route.)

The analysis of a routing algorithm, particularly in a broadcast, and thus mobile, network, is a complex task, in that routing is topology - and load-dependent, and involves, with varying degrees of subtlety, all of the system's protocols. Thus, routing considerations are really a synthesis of most elements of the system design, and as such, the measurement of the algorithm involves at times the study of the interaction of the many system protocols.

TRAFFIC DISTRIBUTION AND EFFICIENCY

Given the patterns of input load on the network, the distribution of traffic flow in the net is an indication of the behavior and efficiency of the routing and initialization algorithms. One may detect the concentration of traffic on specific routes creating congestion while alternate routes are not assigned; thus smaller delay routes may have been ignored in favor of the shorter routes provided by the initialization procedure. To obtain the distribution of traffic flow, the following quantities are to be measured.

(a) the total number of packets received and transmitted at each repeater (obtained in the PRU CumStats)

(b) the fraction of time the transceiver is busy (obtained by snapshot statistics, or in the PRU CumStat by regular sampling of the transceiver's state).

Also, the "point-to-point" nature of this routing algorithm, restricts a packet at a given hop to a single repeater as its immediate destination, does not take advantage of the broadcast nature of the channel, in which several neighbors may actually hear the transmission and be capable of relaying the packet. Thus the following quantity is relevant:

(*) In forthcoming CAP versions, point-to-point routing algorithms will be implemented. As this is done, new experiments will be designed to examine these protocols. For the time being, we assume hierarchical routing.

(c) the number of packets correctly received and discarded because they are destined to other components in the net (obtained in the PRU CumStat).

Moreover, to measure the potential of each neighboring repeater (say, repeater "n") as an immediate destination, it is essential to know the probability of success P(n) repeater n has to correctly receive a broadcast packet. This we do by maintaining in each PRU a table counting the number of successfully received packets from each immediate neighbor. The ratio of the number of packets correctly received from a given neighbor, to the number of packets transmitted by that neighbor, is a measure of P(n).

Most of the above concerns examining the efficiency of operational user networks by observing the actual traffic distribution determined by the routing algorithm. In this experimental net we intend to examine the relationship between traffic distribution and network efficiency. This experiment is aimed at learning whether it is more efficient to route a given throughput on a single branch or to equally split the traffic along the two branchs.

Area of Interest: Traffic distribution and efficiency

See SSI testbed figure(s): a) 1; b) 3,5 [modified]

a)

b)

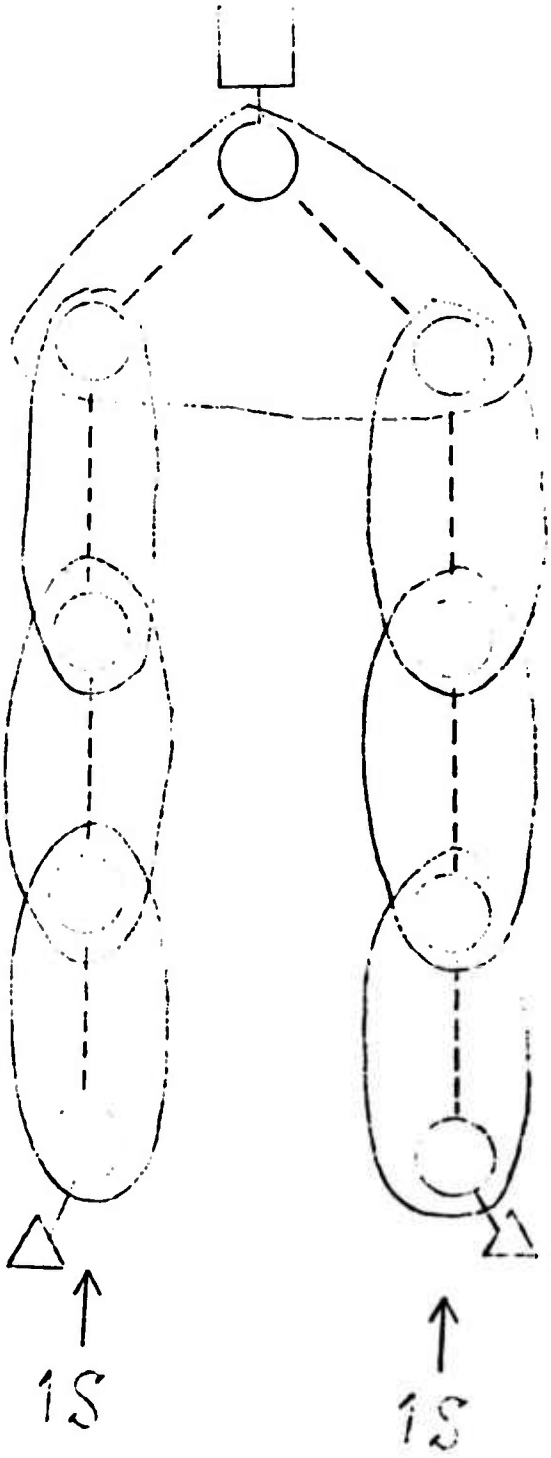

Area of Interest: Traffic distribution and efficiency

Configurations: Two independent paths of repeaters, meeting at the station.

Tools: Cumstats (end-device and PRU)

Parameters: For a range of input rates (minimal channel usage, through saturation), S, apply the traffic as indicated in the accompanying configuration figure.

Data Items: 1) Maximum achievable thourghput (CS-PRU)
2) ETE delay [CS-end-device]

Processing: Compute delay vs throughput for the two methods under the various traffic loads.

ROUTING SENSITIVITY TO TRAFFIC FLOW

The station routing algorithm is intended to speed the flow of traffic by detecting and responding to non-optimal traffic distribution patterns. In this experiment we will use a configuration in which we systematically create unbalanced traffic patterns to observe the station's response and to determine the threshold at which it will respond.

This configuration also allows for a pattern in which the station's response can only result in the shifting of some traffic from one route to another, thereby recreating a mirror image of the original unbalanced traffic pattern, for which the station's only response could be to return the routes to their original pattern again, etc. An example would be where the respective input rates of the three terminals (see figure with configuration) are:

T1: 1s
T2: 5s
T3: 1s

Area of Interest: Routing sensitivity to traffic flow

See SRI testbed figure(s): 6

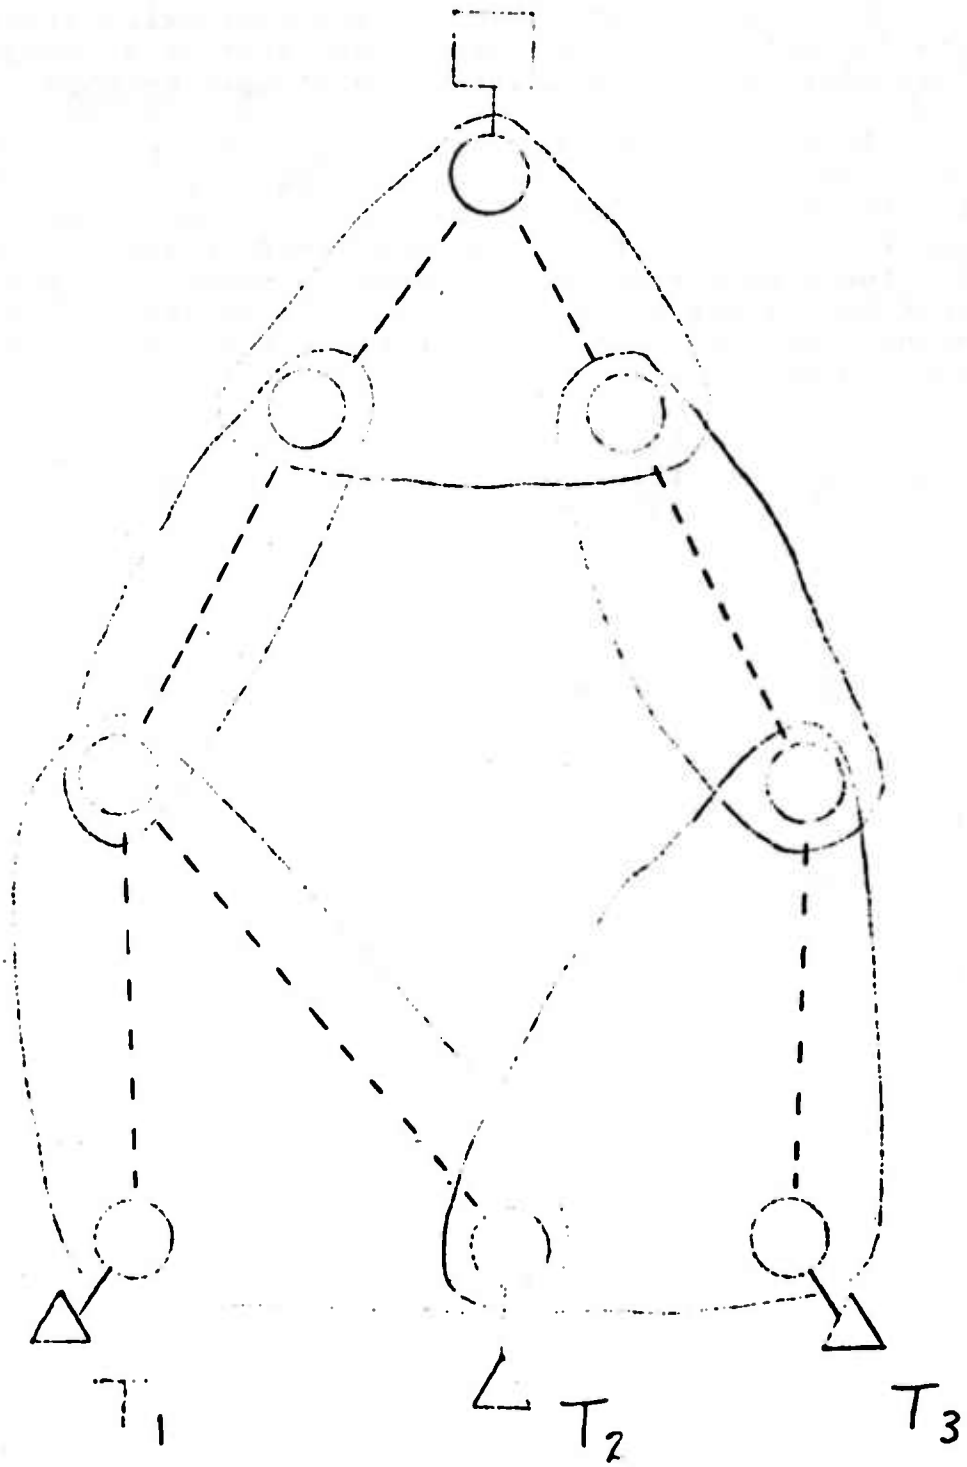

Area of Interest: outing sensitivity to traffic flow

Configurations: Terminals T1 and T3 have separate paths to the station, and terminal T2 can route traffic along either of those paths. The relative input traffic of the three terminals will determine the congestion on each of the two routes.

Tools: Cumstats (end-device and PRU), Snapshots (station)

Data Items:
1) # of packets rx'd and tx'd at each node [CS-PRU]
2) # of packets rx'd at station from each terminal [CS-station] (throughput)
3) route assignments [snapshots-station]
4) ETE delay [CS-end-device]

Parameters: Vary the traffic load on Terminals T1, T2 and T3.

Processing:
1) Find threshold values for routing changes
2) Find effectiveness of routing changes (thoughput, delay)

Comment: In case efficiency is enhanced by uniform distribution of traffic over available branches, we must check for "flipping", i.e. a condition in which the station continually alters T2's route from one branch to the other in an attempt to equalize the traffic flow.

NETWORK INITIALIZATION TIME

This experiment is aimed at the determination of the relationship between network size and initialization time. Using time-stamps of updates in the connectivity and routing tables (initially empty) at the station, the level-by-level progress of initialization can be measured, and the relationship between initialization time and the number of levels and elements per level can be described. The extrapolation of this function to large-scale networks will be of particular interest.

Area of Interest: Network Initialization Time

See SRI testbed figure(s): 1 (STRING)

STRING:

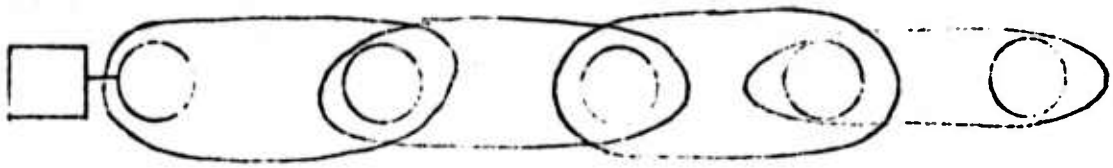

RING:

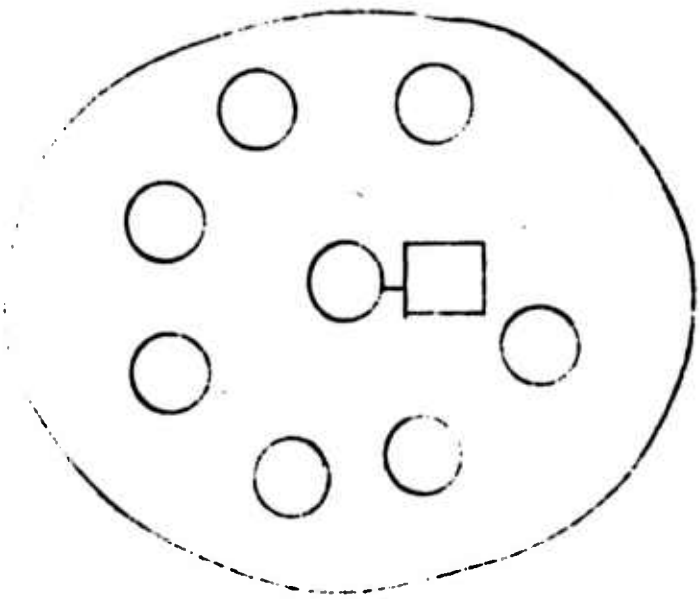

Area of Interest: Network Initialization Time

Configurations: We will first examine simple configurations such as a string (which corresponds to a tree of K levels, with 1 repeater at each level) and a ring (which corresponds to a tree of 1 level and N elements in that level). It will also be interesting to consider the "largest" network that the test-bed can support, in terms of the number of levels and repeaters at each level.

Tools: Station connectivity table
Station routing table

Parameters: K (number of levels), N (number of repeaters/level); ROP frequencies.

Data Items: Station time-stamp of additions to connectivity and routing tables.

Data Reduction: Progress of initialization: component initialization time in relation to K and N.

COMPONENT INITIALIZATION TIME

This experiment is a companion experiment to the Network Initialization Time experiment, with the exception that we consider here the labeling of a component thrown into an already initialized network. The focus here is on the time of occurrence of various events (first tx of unlabeled ROP; reception of label packet) from the point of view of the added component itself.

Area of Interest: Component Initialization Time

See SRI testbed figure(s): To Be Arranged

a)

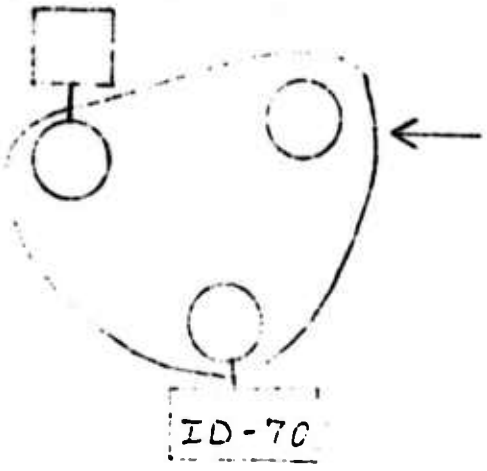

b)

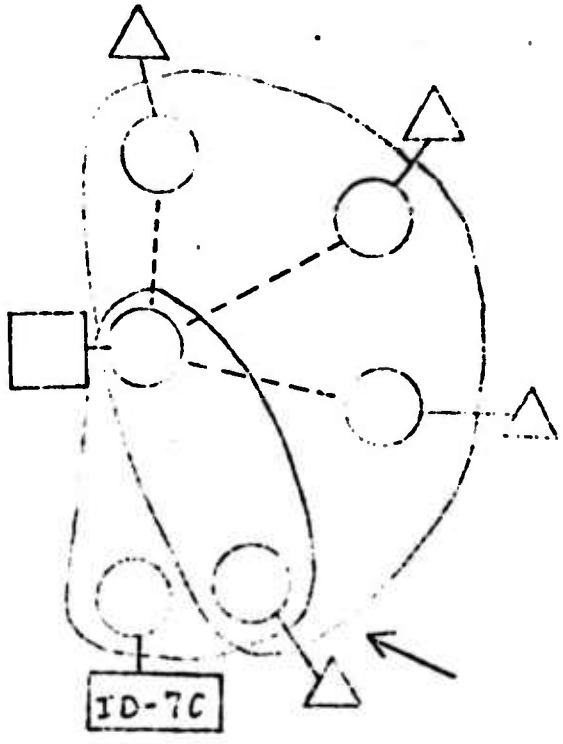

c)

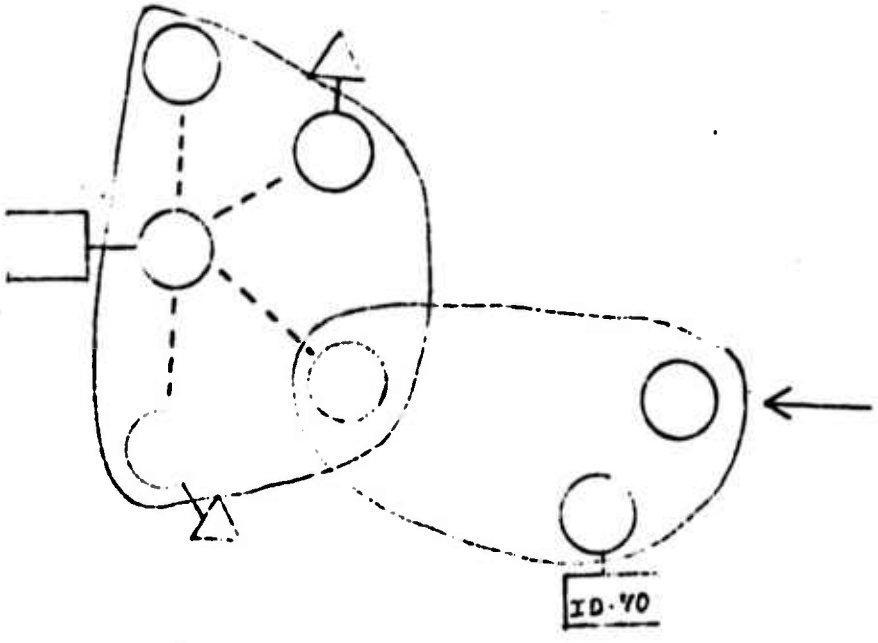

Area of Interest: Component Initialization Time

Configurations: Three examples are selected. The first is the simplest: the network consists of only a station; the second is a one-level network, to which is added the unlabeled PRU at the same level (here, the station has to process a larger number of ROP's relayed by the other components); the third involves a multihop network. In the figure, the arrow points to the component to be initialized.

Tools: Interdata-70 with superhose monitor

Parameters: varying traffic rates from pool of terminals, varying ROP frequencies.

Data Items: ID-70 trace of ROPs, label packet(s), labeled ROPs.

Data Reduction: Delay from 1st ROP to rx of label packet.

RESPONSE TO COMPONENT FAILURE

A key consideration in network design is the ability to react dynamically, swiftly and efficiently to the sudden loss of a component. In the PRNET we are particluarly concerned with the response of the network to the unexpected failure of a repeater. We want to know how quickly the controlling station can detect and respond to the failure, and also the effects on local traffic of the failure and of the station's response.

The measures of interest will be the lapse of time between failure and the station's response (measured as the relabeling of components affected by the failure); packet delays at affected terminals before failure, after failure and prior to the station's response, and following the station's response. These local effects (around the failed component) will be measured in terms of local single hop delays, and one-way delays (time from source to destination) obtained via the Pickup Packets.

Area of Interest: Response to Component Failure

See SRI testbed figure(s): Possibility: 5 [illegible]

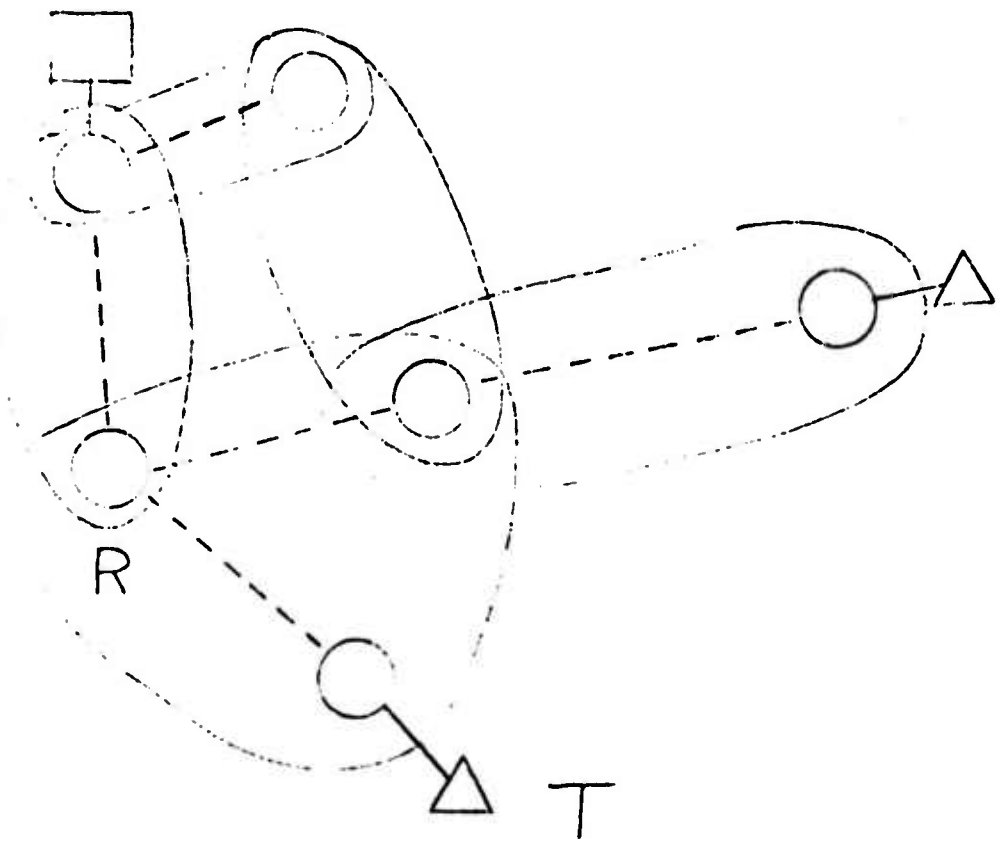

Area of Interest: Response to Component Failure

Configurations: We consider the configuration depicted, in which repeater R is the failing repeater. Terminal T is the directly affected component.

Tools: Pickup Packets; Station labels

Parameters: Pickup packets generated at given rates from both terminals.

Data Items: Time-stamps, routes taken [PP]; time new label generated [station], local single-hop delay; number of alt-routed packets between failure and station response.

Data Reduction: On a time-line, show delay between component failure and terminal relabeling, with intermediate events (# of alt-rt pkts rx'd), and plot delay (single-hop and one-way) [PP] as shown below:

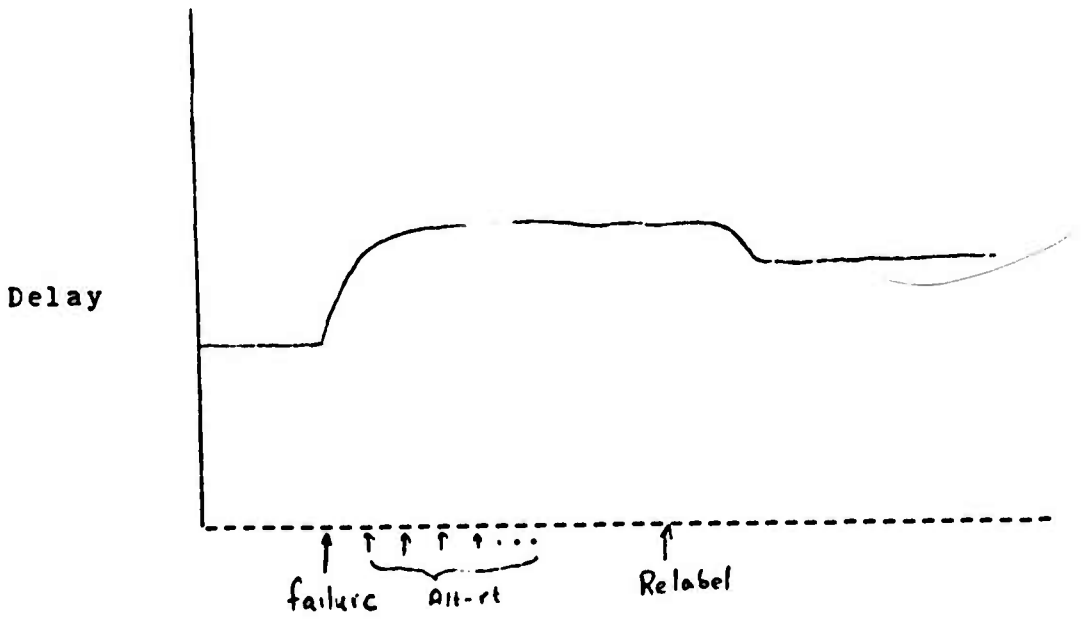

MOBILE TERMINAL

With mobile terminals, the issue is the efficient relaying of traffic originating at the terminals as they move from one repeater's range to another's. The measures of efficiency are the percent of time the terminal PRU is using the wrong label (and thus is alternately routing), and the single-hop delay from the terminal to its relaying repeater, and ETE delay.

Area of Interest: Mobile Terminals

See SRI testbed figure(s): Possibility: 5 (modified)

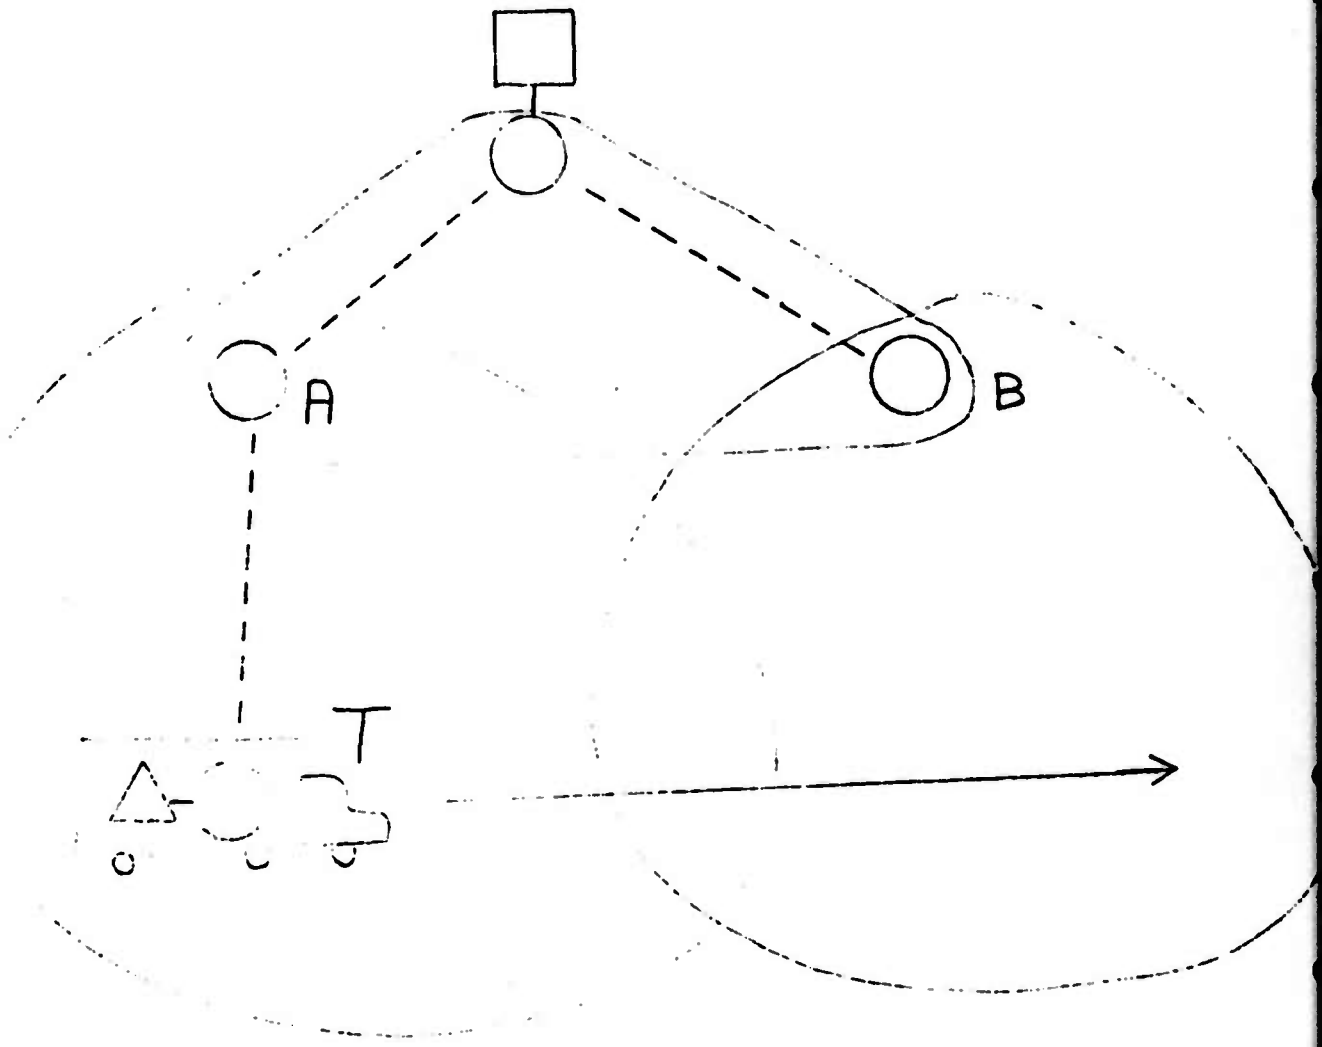

Area of Interest: Mobile Terminals

Configuration: A simple configuration for this experiment has been selected. As terminal T moves, it leaves repeater A's range to enter repeater B's range.

Tools: CumStat (PRU and End-Device); Pickup Packets

Parameters: ROP rates, # of tx's before alt-routing.

Data Items: Histogrms of tx's to success [CS-PRU]
ETE delay [CS-End Device]
1-hop delay [PP]

Processing: Use alt-routing data from CumStat histograms to compute the percent of time the terminal PRU was using the wrong label. Correlate with the ETE and 1-hop delays.

EFFECTS OF 'N' ON CSMA

The analysis of a simple two-hop network configuration using slotted ALOHA and CSMA was published in PRTNs 247, 248, and 249. In this experiment we intend to use the same configuration and attempt to verify some of the analytic results, mostly those relating to CSMA (since it is implemented in the PRNET). Basically the analysis has considered the fully-connected two-hop network configuration in which all repeaters are within range and in line-of-sight of each other and of the station. With each repeater is associated a population of terminals generating traffic which is destined to the station. Terminals and repeaters follow the nonpersistent CSMA protocol.

Area of Interest: Effect of # of Repeaters on Access Mode Performance

See SFI testbed figure(s): To Be Prepared

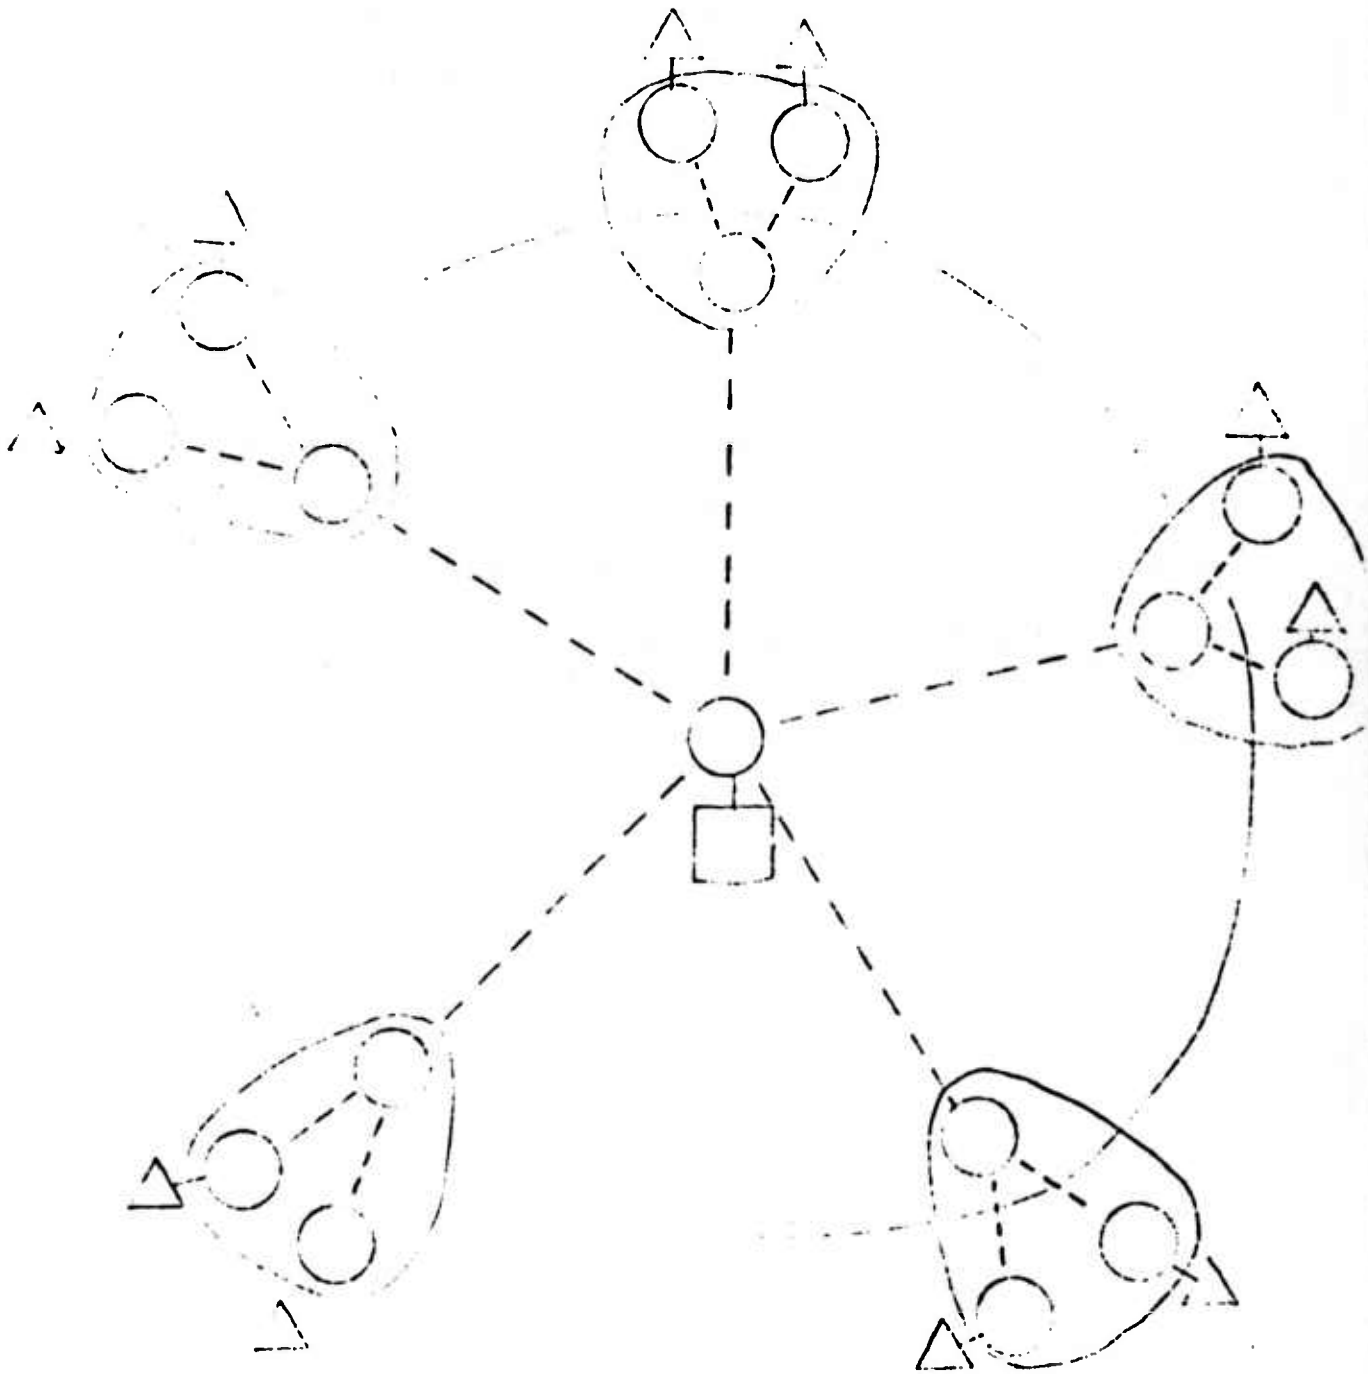

Area of Interest: Effect of # of Repeaters on Access Mode Performance

Configuration: A ring of repeaters around the station, each repeater with it's population of terminals that have connectivity only with their repeater.

Tools: CumStats (end-device and PRU), Pickup packets

Parameters: Traffic sources sending information and PPs (of various sizes) at various rates s such that
0 LE Ns LE f(N) (as per analysis)
e.g. if N = 10 then analysis indicates f(10) = 0.43
Other parameters: (i) retransmission delay;
(ii) Initial transmission delay (0, i.e. immediate first transmission IFT; and non-zero, i.e. delayed first trnasmission DFT).

Data Items: # of tx's to success [CS-PRU]
# of successful tx's [CS-PRU]
Round trip delay [CS-end device]
Access delay (terminal to repeater) [PP]
Delay at repeater [PP]

Processing: Mean and standard deviation of items above

Delay vs throughput for various N (0 < N <= 10)
Maximum throughput vs N

Comments: Do same with ALOHA schemes

PART III

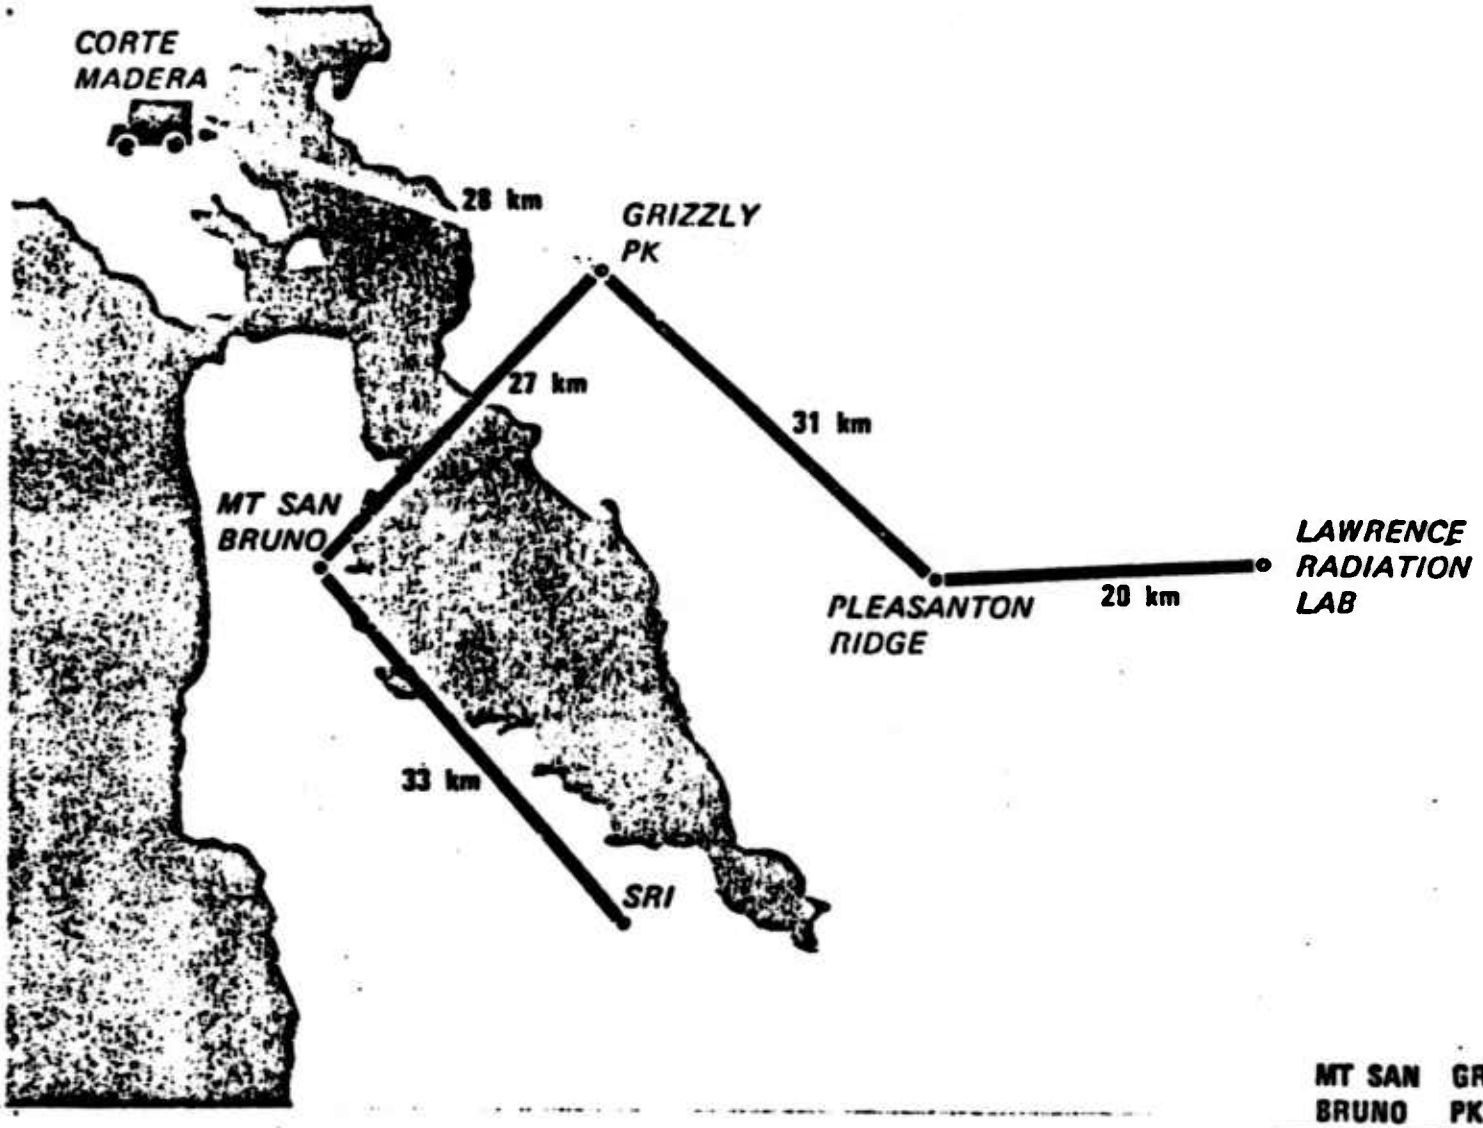

10 km
10 mi

6 PACKET RADIO UNITS
15 POSSIBLE PATHS
5 PATHS PROVIDED
10 PATHS DENIED

5 PACKET RADIO UNITS
10 POTENTIAL PATHS
4 PATHS PROVIDED
6 PATHS DENIED

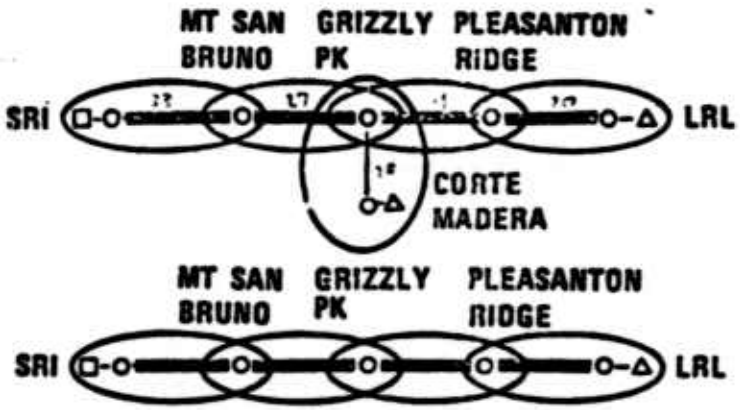

Figure 1

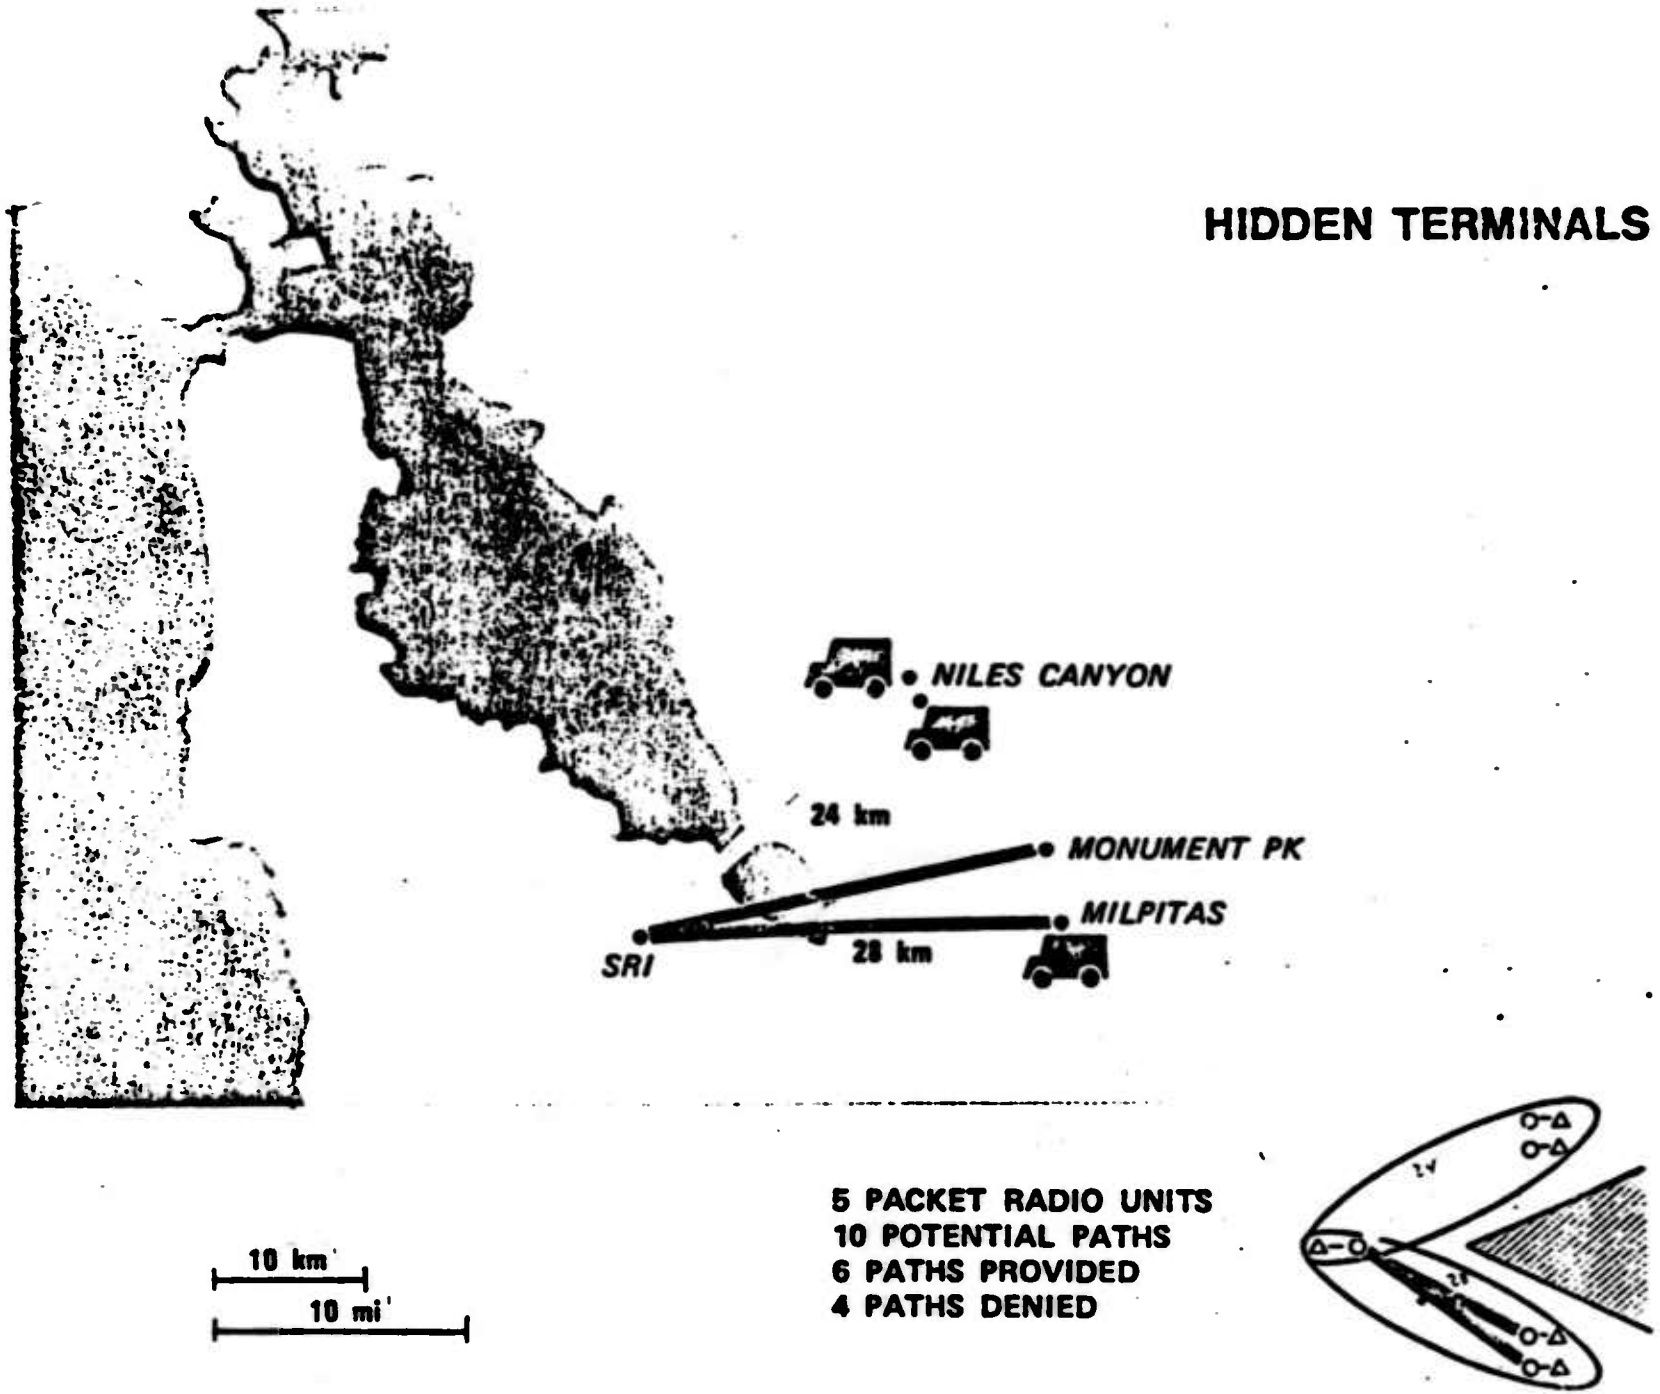
HIDDEN TERMINALS
NILES CANYON
24 km
MONUMENT PK
MILPITAS
SRI
28 km
10 km
10 mi
5 PACKET RADIO UNITS
10 POTENTIAL PATHS
6 PATHS PROVIDED
4 PATHS DENIED

Figure 2

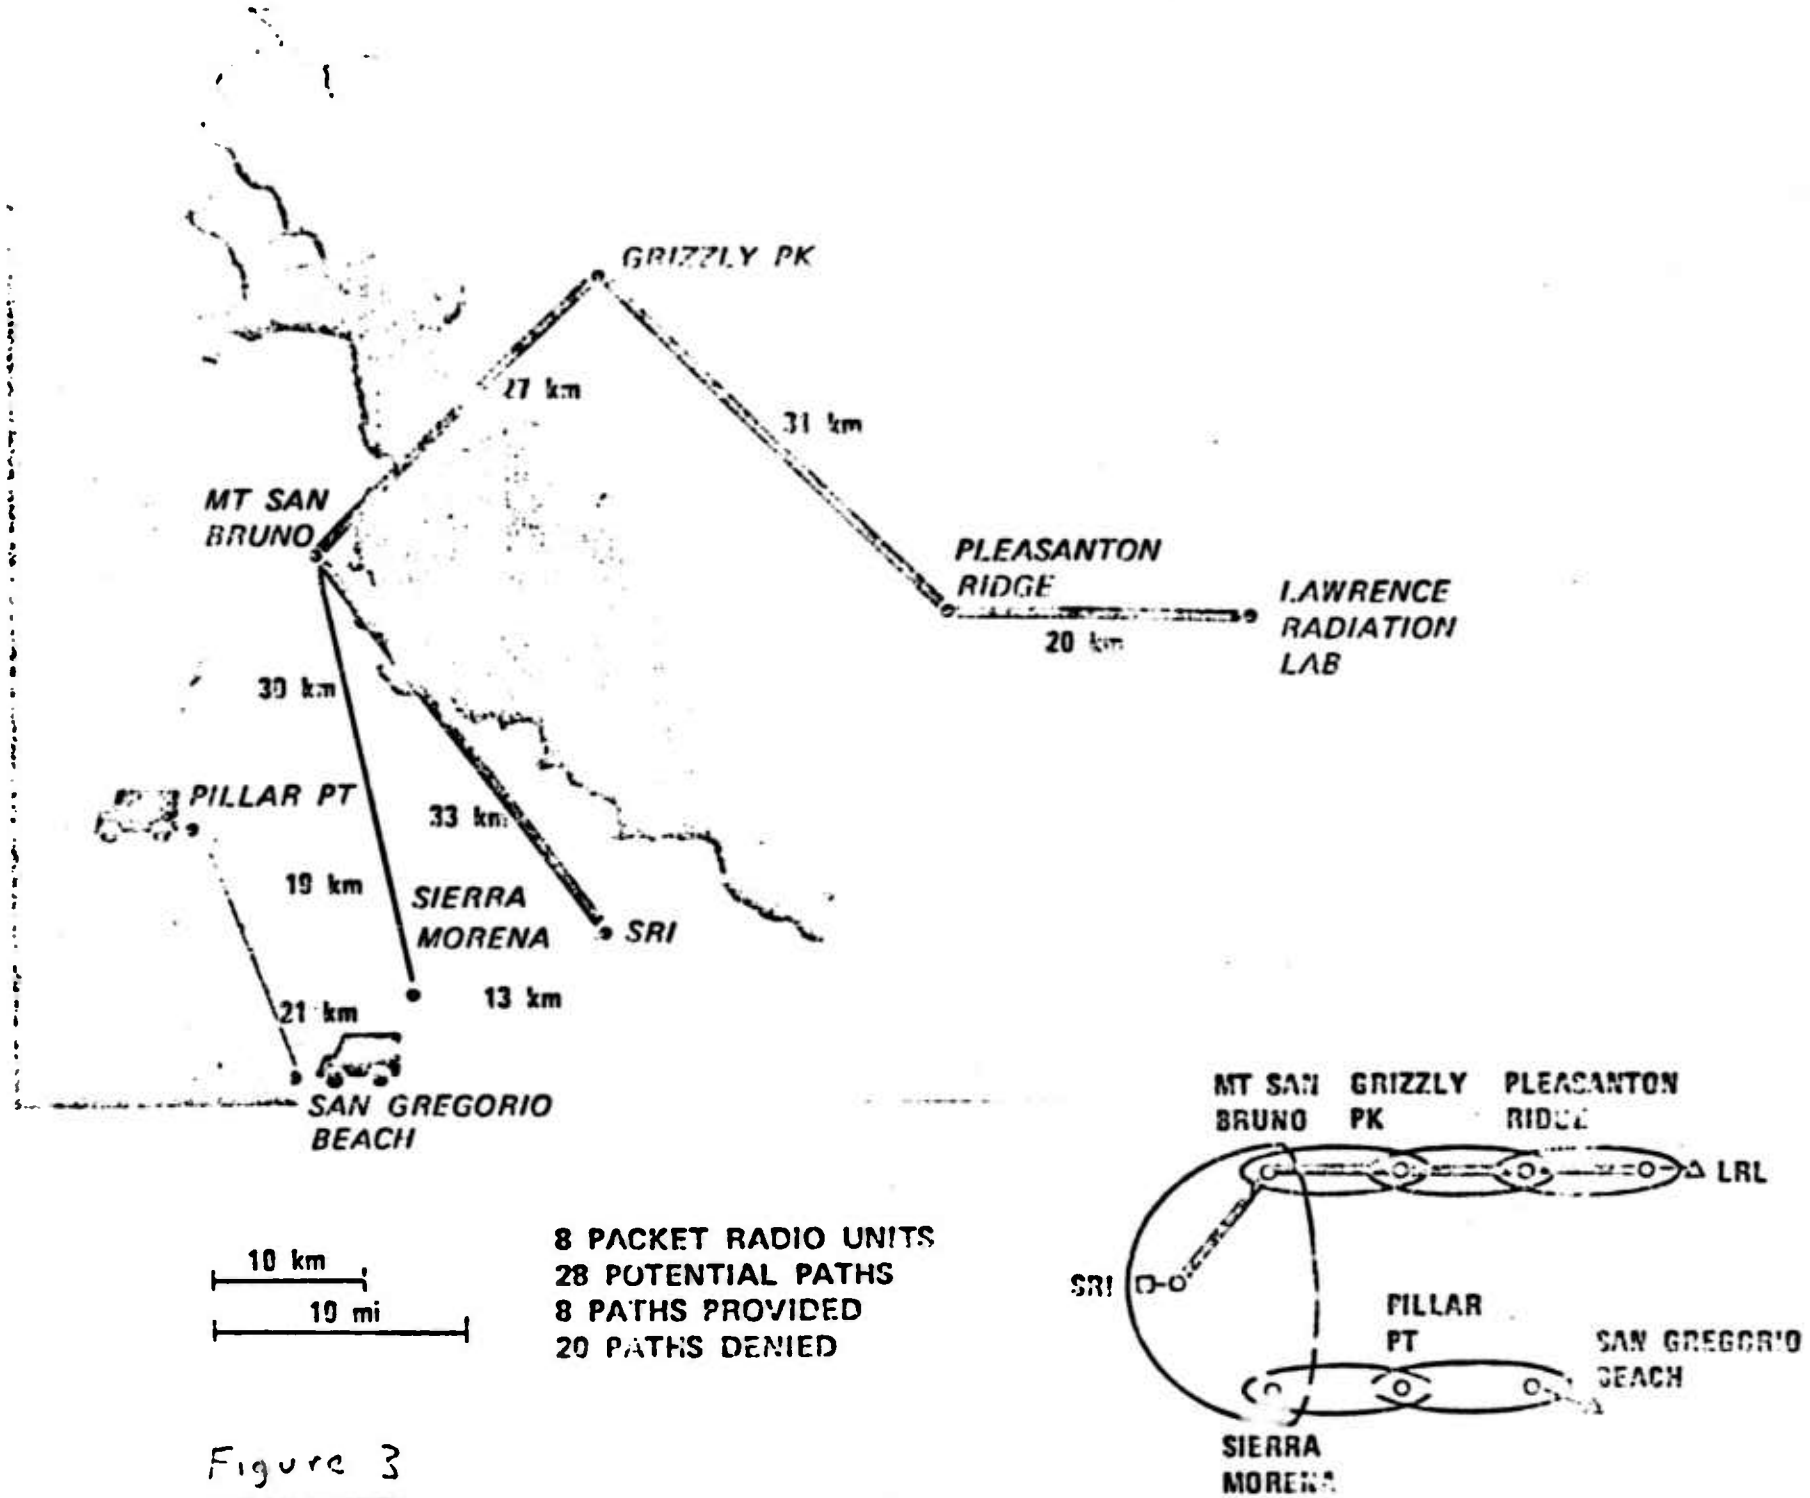
GRIZZLY PK
27 km
31 km
MT SAN BRUNO
PLEASANTON RIDGE
LAWRENCE RADIATION LAB
20 km
30 km
PILLAR PT
33 km
19 km
SIERRA MORENA
SRI
13 km
21 km
SAN GREGORIO BEACH
10 km
10 mi
8 PACKET RADIO UNITS
28 POTENTIAL PATHS
8 PATHS PROVIDED
20 PATHS DENIED
MT SAN BRUNO
GRIZZLY PK
PLEASANTON RIDGE
LRL
SRI
PILLAR PT
SAN GREGORIO BEACH
SIERRA MORENA

Figure 3

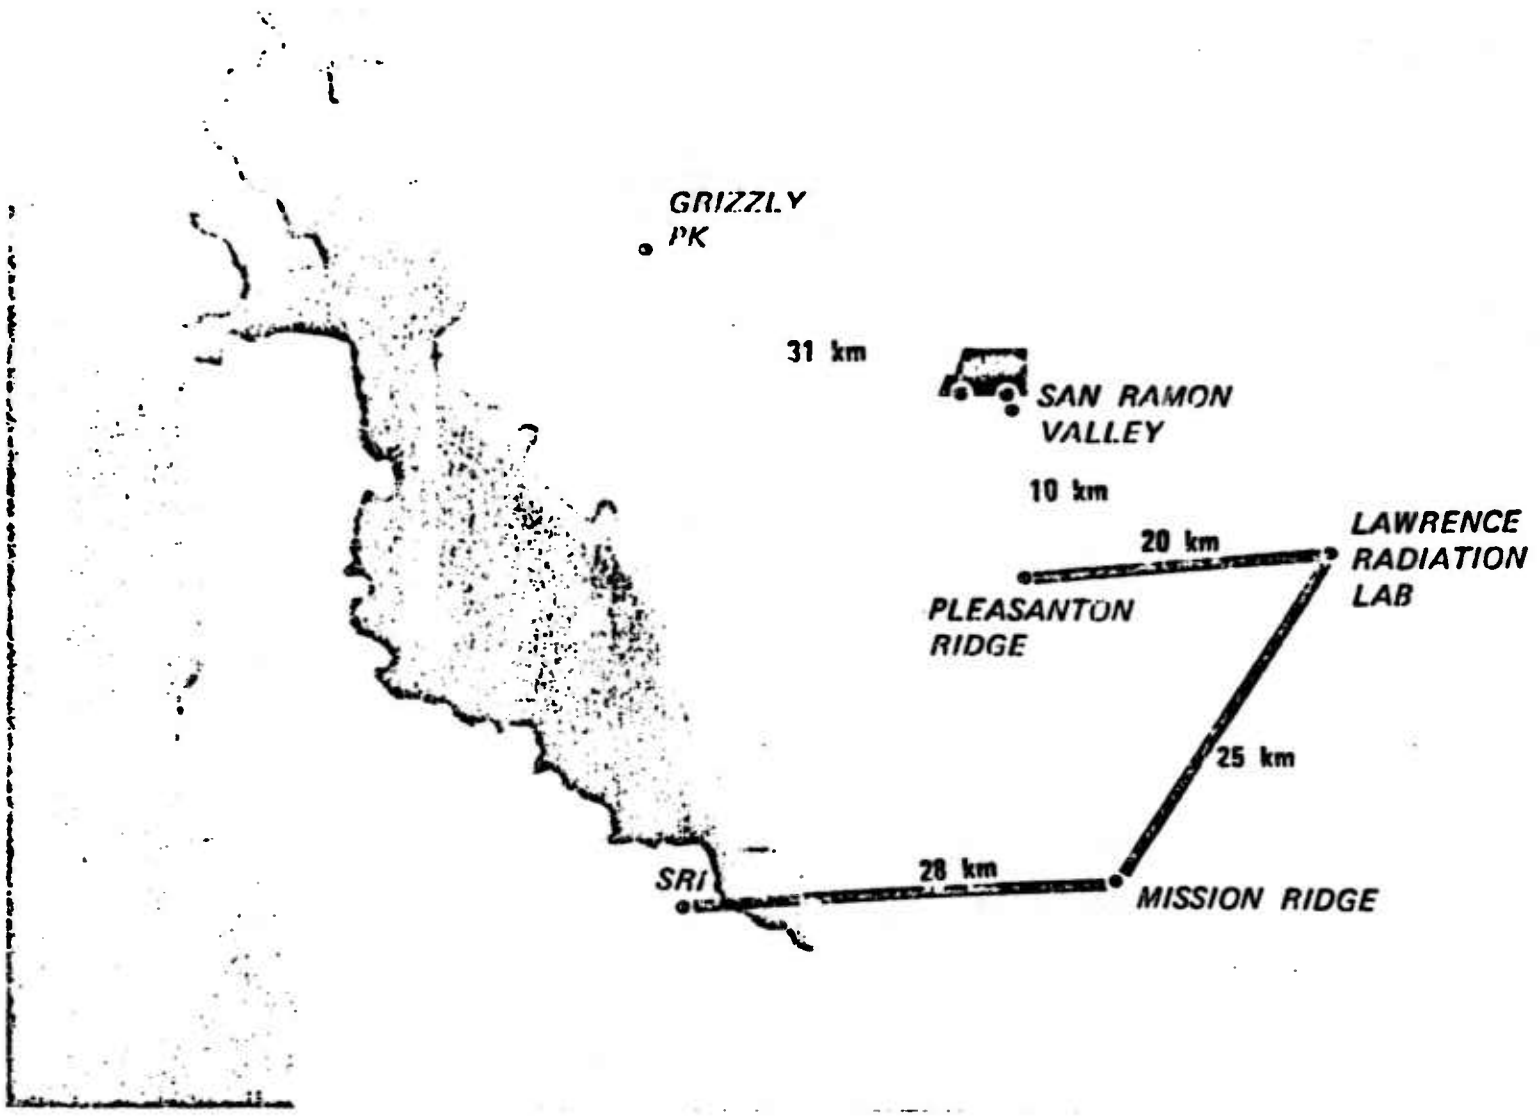

10 km
10 mi

6 PACKET RADIO UNITS
15 POTENTIAL PATHS
5 PATHS PROVIDED
10 PATHS DENIED

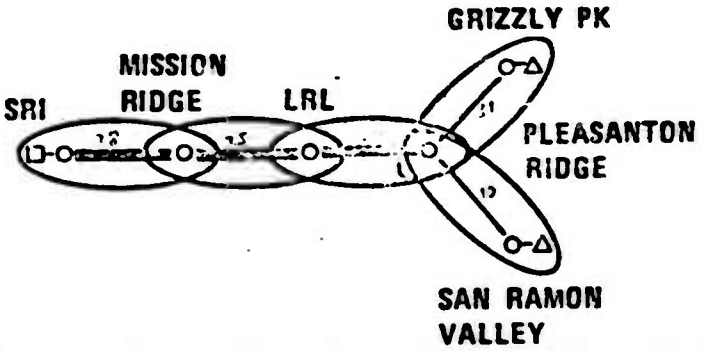

Figure 4

THE SRI AREA MINI NET

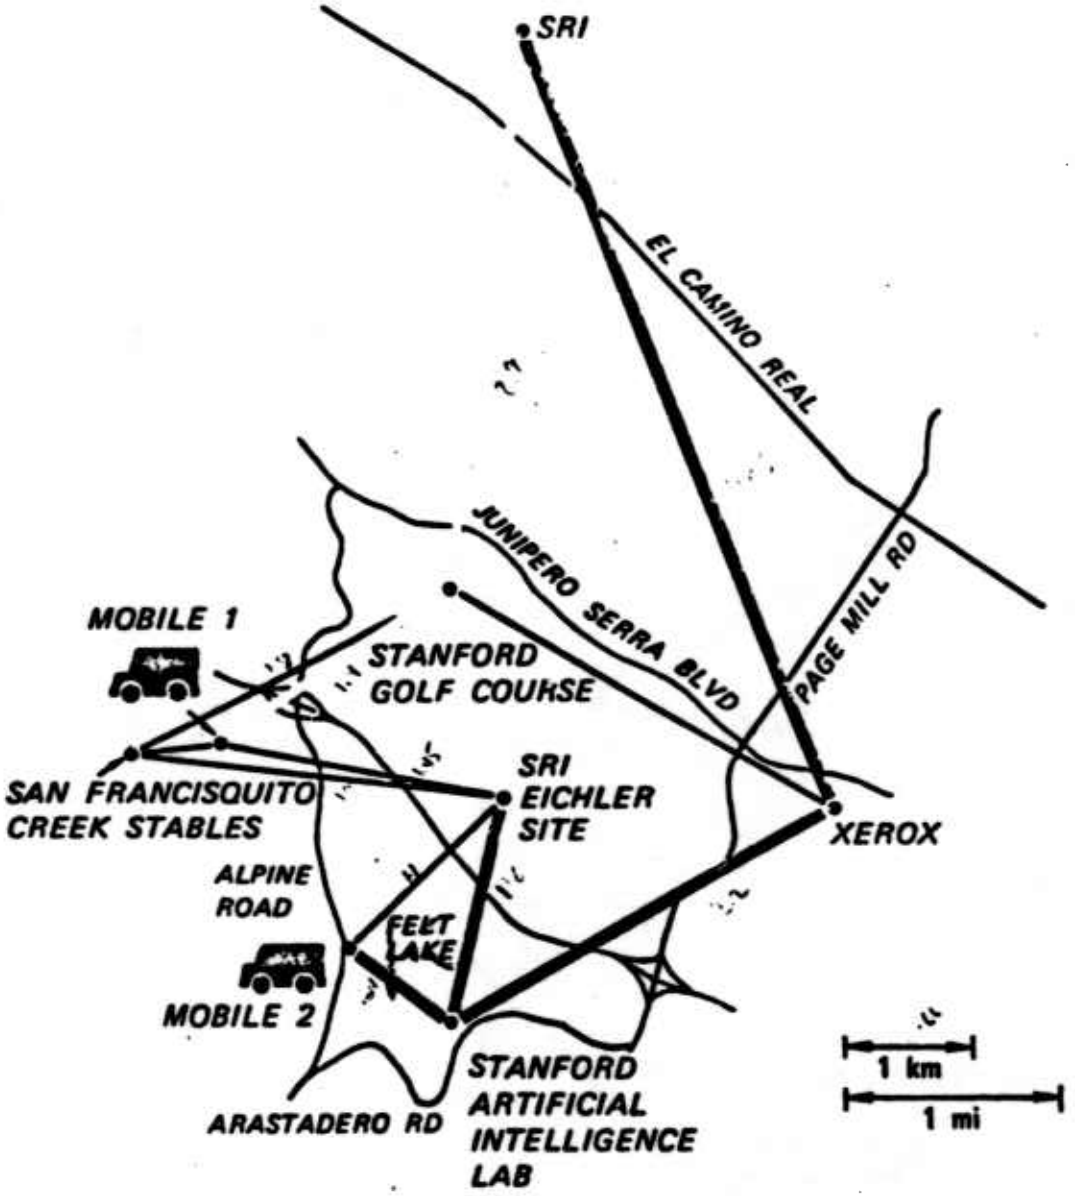

8 PACKET RADIO UNITS
28 PATHS POTENTIAL PATHS
11 PATHS PROVIDED
17 PATHS DENIED

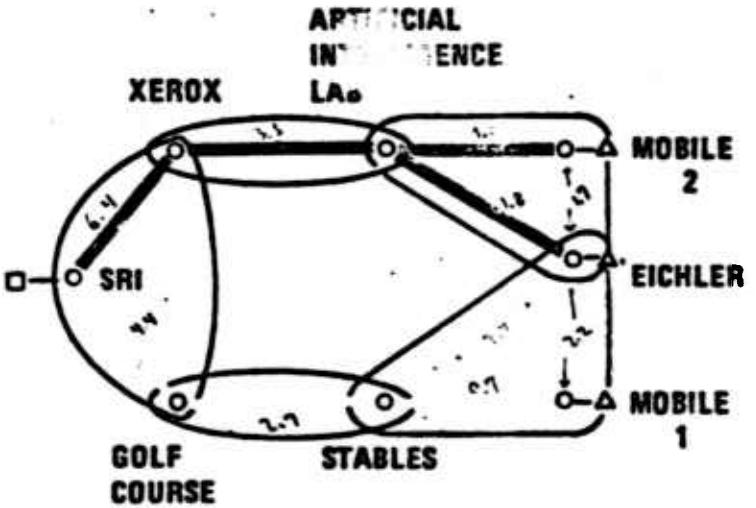

Figure 5

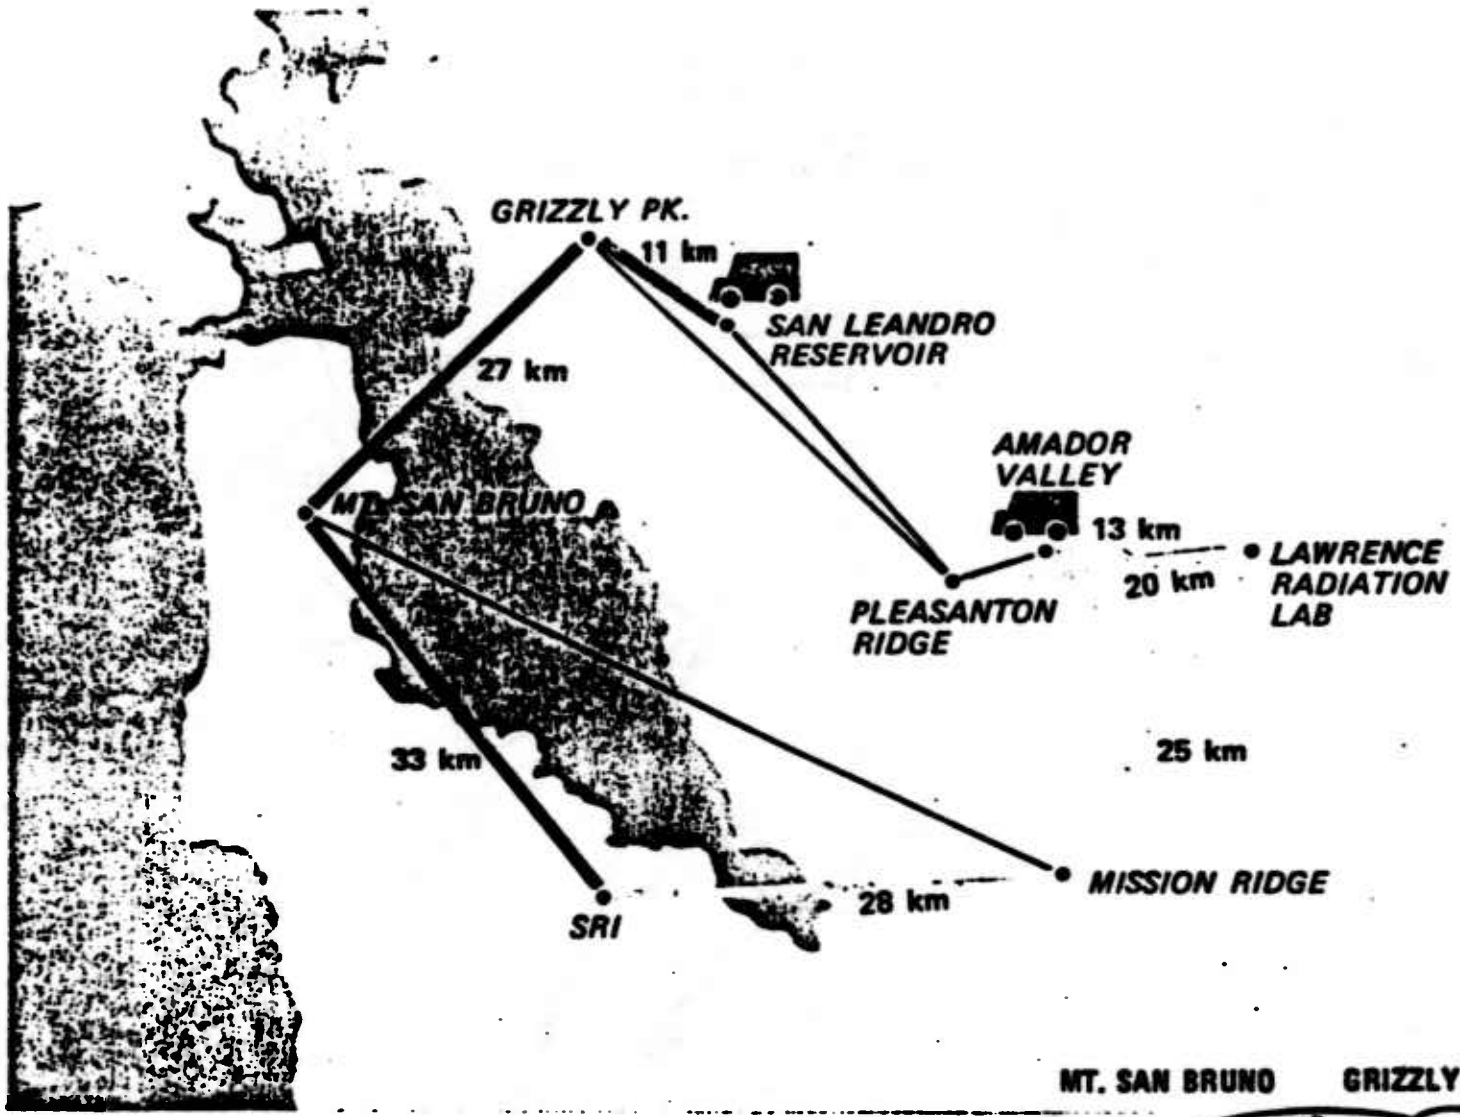

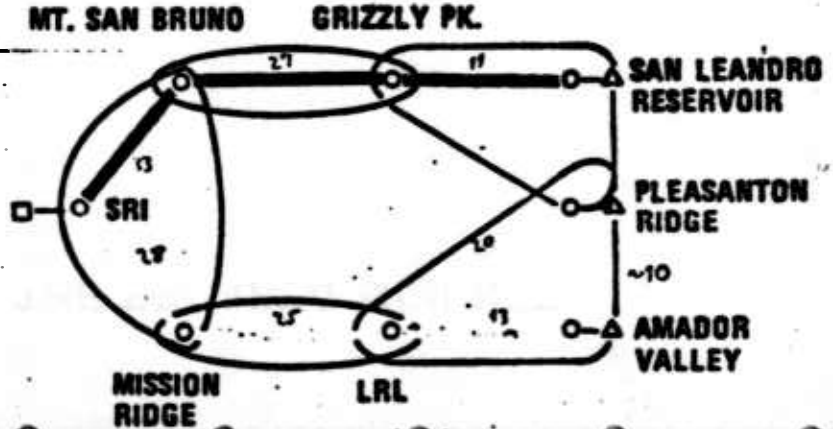

10 km
10 mi

8 PACKET RADIO UNITS
28 POTENTIAL PATHS
11 PATHS PROVIDED
17 PATHS DENIED

Figure 6

PART IV

Tentative Working Schedule

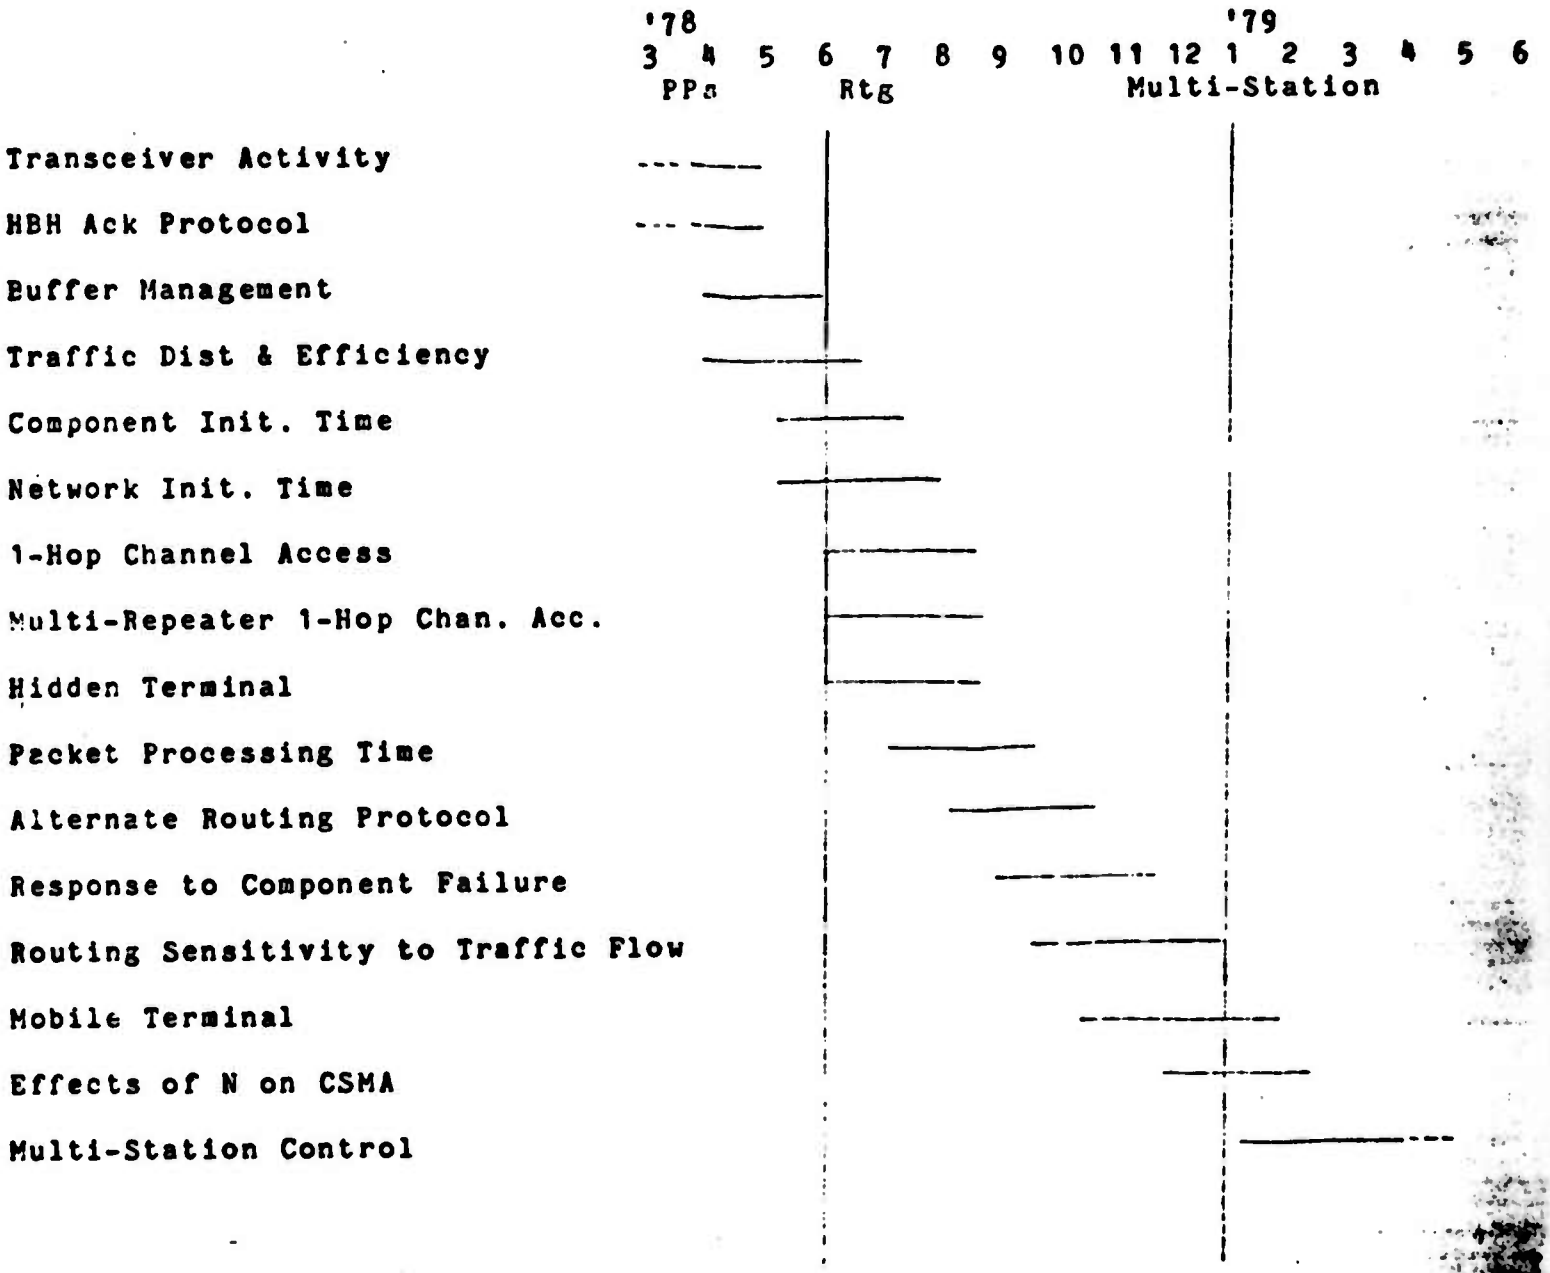

S